

Timed Soft Concurrent Constraint Programs: An Interleaved and a Parallel Approach

Stefano Bistarelli

*Dipartimento di Matematica e Informatica, Università di Perugia
Via Vanvitelli 1, 06123 Perugia Italy
E-mail: vista@dmf.unipg.it*

Maurizio Gabbrielli

*Dipartimento di Scienze dell'Informazione, Università di Bologna
Via Zamboni 33, 40126 Bologna, Italy
E-mail: gabbri@cs.unibo.it*

Maria Chiara Meo

*Dipartimento di Economia, Università "G. D'Annunzio"
Viale Pindaro 42, 65127 Pescara, Italy
E-mail: cmeo@unich.it*

Francesco Santini

*Centrum Wiskunde & Informatica (CWI)
Science Park 123, 1098XG Amsterdam, The Netherlands
E-mail: F.Santini@cwi.nl*

submitted 20 May 2012; revised 15 February 2014; accepted 18 February 2014

Abstract

We propose a timed and soft extension of Concurrent Constraint Programming. The time extension is based on the hypothesis of *bounded asynchrony*: the computation takes a bounded period of time and is measured by a discrete global clock. Action prefixing is then considered as the syntactic marker which distinguishes a time instant from the next one. Supported by soft constraints instead of crisp ones, *tell* and *ask* agents are now equipped with a preference (or consistency) threshold which is used to determine their success or suspension. In the paper we provide a language to describe the agents behavior, together with its operational and denotational semantics, for which we also prove the compositionality and correctness properties. After presenting a semantics using maximal parallelism of actions, we also describe a version for their interleaving on a single processor (with maximal parallelism for time elapsing). Coordinating agents that need to take decisions both on preference values and time events may benefit from this language. To appear in *Theory and Practice of Logic Programming (TPLP)*.

KEYWORDS: Soft Concurrent Constraint Programming, Timed Concurrent Constraint Programming, Interleaving, Parallelism.

1 Introduction

Time is a particularly important aspect of cooperative environments. In many “real-life” computer applications, the activities have a temporal duration (that can be even interrupted) and the coordination of such activities has to take into consideration this timeliness property. The interacting actors are mutually influenced by their actions, meaning that A reacts accordingly to the timing and quantitative aspects related to B ’s behavior, and vice versa. In fact, these interactions can be often related to quantities to be measured or minimized/maximized, in order to take actions depending from these scores: consider, for example, some generic communicating agents that need to take decisions on a (monetary) cost or a (fuzzy) preference for a shared resource. They both need to coordinate through time-dependent and preference-based decisions.

A practical example of such agents corresponds, for example, to software agents that need to negotiate some service-level agreement on a resource, or a service, with time-related side-conditions. For instance, a fitting example is given by auction schemes, where the seller/bidder agents need to agree on a preference for a given prize (e.g., a monetary cost). At the same time, the agents have to respect some timeout and alarm events, respectively representing the absence and the presence of bids for the prize (for instance). The language we present in this paper is well suited for this kind of interactions, as Section 5 shows with examples.

The *Timed Concurrent Constraint Programming (tccp)*, a timed extension of the pure formalism of *Concurrent Constraint Programming (ccp)* (Saraswat 1989), has been introduced in (de Boer et al. 2000). The language is based on the hypothesis of *bounded asynchrony* (Saraswat et al. 1996): computation takes a bounded period of time rather than being instantaneous as in the concurrent synchronous languages **ESTEREL** (Berry and Gonthier 1992), **LUSTRE** (Halbwachs et al. 1991), **SIGNAL** (le Guernic et al. 1991) and **Statecharts** (Harel 1987). Time itself is measured by a discrete global clock, i.e., the internal clock of the *tccp* process. In (de Boer et al. 2000) the authors also introduced *timed reactive sequences*, which describe the reaction of a *tccp* process to the input of the external environment, at each moment in time. Formally, such a reaction is a pair of constraints $\langle c, d \rangle$, where c is the input and d is the constraint produced by the process in response to c .

Soft constraints (Bistarelli 2004; Bistarelli et al. 1997) extend classical constraints to represent multiple consistency levels, and thus provide a way to express preferences, fuzziness, and uncertainty. The *ccp* framework has been extended to work with soft constraints (Bistarelli et al. 2006), and the resulting framework is named *Soft Concurrent Constraint Programming (sccp)*. With respect to *ccp*, in *sccp* the *tell* and *ask* agents are equipped with a preference (or consistency) threshold, which is used to determine their success, failure, or suspension, as well as to prune the search; these preferences should preferably be satisfied but not necessarily (i.e. over-constrained problems). We adopt soft constraints instead of crisp ones, since classic constraints show evident limitations when trying to represent real-life scenarios, where the knowledge is not completely available nor crisp.

In this paper, we introduce a timed and soft extension of *ccp* that we call *Timed*

Soft Concurrent Constraint Programming (*tsccp*), inheriting from both *tccp* and *sccp* at the same time. In *tsccp*, we directly introduce a timed interpretation of the usual programming constructs of *sccp*, by identifying a time-unit with the time needed for the execution of a basic *sccp* action (ask and tell), and by interpreting action prefixing as the next-time operator. An explicit timing primitive is also introduced in order to allow for the specification of timeouts. In the first place, the parallel operator of *tsccp* is first interpreted in terms of maximal parallelism, as in (de Boer et al. 2000). Secondly, we also consider a different paradigm, where the parallel operator is interpreted in terms of interleaving, however assuming maximal parallelism for actions depending on time. In other words, time passes for all the parallel processes involved in a computation. This approach, analogous to that one adopted in (de Boer et al. 2004), is different from that one of (de Boer et al. 2000; Bistarelli et al. 2008) (where maximal parallelism was assumed for any kind of action), and it is also different from the one considered in (Busi et al. 2000), where time does not elapse for timeout constructs. This can be accomplished by allowing all the time-only dependent actions (τ -transitions) to concurrently run with at most one action manipulating the store (a ω -transition).

The paper extends the results in (Bistarelli et al. 2008) by providing new semantics that allows maximal parallelism for time elapsing and an interleaving model for basic computation steps (see Section 7). This new language is called *tsccp with interleaving*, i.e., *tsccp-i*, to distinguish it from the version allowing maximal parallelism of all actions. According to the maximal parallelism policy (applied, for example, in the original works as (Saraswat 1989) and (Saraswat et al. 1994)), at each moment every enabled agent of the system is activated, while in the interleaving paradigm only one of the enabled agents is executed instead. This second paradigm is more realistic if we consider limited resources, since it does not imply the existence of an unbounded number of processors. However, in (de Boer et al. 2000) it is shown that the notion of maximal parallelism of *tsccp* is more expressive than the notion of interleaving parallelism of other concurrent constraint languages. The presence of maximal parallelism can force the computation to discard some (non-enabled) branches which could become enabled later on (because of the information produced by parallel agents), while this is not possible when considering an interleaving model. Therefore, *tsccp* is sensitive to delays in adding constraints to the store, whereas this is not the case for *ccp* and *tsccp-i*.

The rest of the paper is organized as follows: in Section 2 we summarize the most important background notions and frameworks from which *tsccp* derives, i.e. *tccp* and *sccp*. In Section 3 we present the *tsccp* language, and in Section 4 describes the operational semantics of *tscc* agents. Section 5 better explains the programming idioms as *timeout* and *interrupt*, exemplifies the use of timed paradigms in the *tscc* language and shows an application example on modeling an auction interaction among several bidders and a single auctioneer. Section 6 describes the denotational semantics for *tsccp*, and proves the denotational model correctness with the aid of *connected reactive sequences*. Section 7 explains the semantics for interleaving with maximal parallelism of time-elapsing actions (i.e. the *tsccp-i* language), while Section 8 describes a timeline for the execution of three parallel agents in *tsccp-i*.

Section 9 describes the denotational semantics of *tscap-i* and proves the correctness of the denotational model. Section 10 reports the related work and, at last, Section 11 concludes by also indicating future research.

2 Background

2.1 Soft Constraints

A *soft constraint* (Bistarelli et al. 1997; Bistarelli 2004) may be seen as a constraint where each instantiation of its variables has an associated value from a partially ordered set which can be interpreted as a set of preference values. Combining constraints will then have to take into account such additional values, and thus the formalism has also to provide suitable operations for combination (\times) and comparison ($+$) of tuples of values and constraints. This is why this formalization is based on the concept of *c-semiring* (Bistarelli et al. 1997; Bistarelli 2004), called just *semiring* in the rest of the paper.

Semirings. A semiring is a tuple $\langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$ such that: *i)* A is a set and $\mathbf{0}, \mathbf{1} \in A$; *ii)* $+$ is commutative, associative and $\mathbf{0}$ is its unit element; *iii)* \times is associative, distributes over $+$, $\mathbf{1}$ is its unit element and $\mathbf{0}$ is its absorbing element. A *c-semiring* is a semiring $\langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$ such that: $+$ is idempotent, $\mathbf{1}$ is its absorbing element and \times is commutative. Let us consider the relation \leq_S over A such that $a \leq_S b$ iff $a + b = b$. Then, it is possible to prove that (see (Bistarelli et al. 1997)): *i)* \leq_S is a partial order; *ii)* $+$ and \times are monotone on \leq_S ; *iii)* $\mathbf{0}$ is its minimum and $\mathbf{1}$ its maximum; *iv)* $\langle A, \leq_S \rangle$ is a complete lattice (a complete lattice is a partially ordered set in which all subsets have both a supremum and an infimum) and, for all $a, b \in A$, $a + b = \text{lub}(a, b)$ (where *lub* is the *least upper bound*).

Moreover, if \times is idempotent, then: $+$ distributes over \times ; $\langle A, \leq_S \rangle$ is a complete distributive lattice and \times its *glb* (*greatest lower bound*). Informally, the relation \leq_S gives us a way to compare semiring values and constraints. In fact, when we have $a \leq_S b$, we will say that *b is better than a*. In the following, when the semiring will be clear from the context, $a \leq_S b$ will be often indicated by $a \leq b$.

Constraint System. Given a semiring $S = \langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$ and an ordered set of variables V over a finite domain D , a *soft constraint* is a function which, given an assignment $\eta : V \rightarrow D$ of the variables, returns a value of the semiring. Using this notation $\mathcal{C} = \eta \rightarrow A$ is the set of all possible constraints that can be built starting from S , D and V .

Any function in \mathcal{C} involves all the variables in V , but we impose that it depends on the assignment of only a finite subset of them. So, for instance, a binary constraint $c_{x,y}$ over variables x and y , is a function $c_{x,y} : (V \rightarrow D) \rightarrow A$, but it depends only on the assignment of variables $\{x, y\} \subseteq V$ (the *support* of the constraint, or *scope*). Note that $c\eta[v := d_1]$ means $c\eta'$ where η' is η modified with the assignment $v := d_1$ (that is the operator $[\]$ has precedence over application). Note also that $c\eta$ is the application of a constraint function $c : (V \rightarrow D) \rightarrow A$ to a function $\eta : V \rightarrow D$; what we obtain, is a semiring value $c\eta$.

The partial order \leq_S over \mathcal{C} can be easily extended among constraints by defining $c_1 \sqsubseteq c_2 \Leftrightarrow c_1\eta \leq c_2\eta$, for each possible η .

Combining and projecting soft constraints. Given the set \mathcal{C} , the combination function $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ is defined as $(c_1 \otimes c_2)\eta = c_1\eta \times c_2\eta$ (see also (Bistarelli et al. 1997; Bistarelli 2004; Bistarelli et al. 2006)). Informally, performing the \otimes between two constraints means building a new constraint whose support involves all the variables of the original ones, and which associates with each tuple of domain values for such variables a semiring element which is obtained by multiplying the elements associated by the original constraints to the appropriate sub-tuples.

Given a constraint $c \in \mathcal{C}$ and a variable $v \in V$, the *projection* (Bistarelli et al. 1997; Bistarelli 2004; Bistarelli et al. 2006) of c over $V - \{v\}$, written $c \Downarrow_{(V - \{v\})}$ is the constraint c' s.t. $c'\eta = \sum_{d \in D} c\eta[v := d]$. Informally, projecting means eliminating some variables from the support. This is done by associating with each tuple over the remaining variables a semiring element which is the sum of the elements associated by the original constraint to all the extensions of this tuple over the eliminated variables.

We define also a function \bar{a} (Bistarelli 2004; Bistarelli et al. 2006) as the function that returns the semiring value a for all assignments η , that is, $\bar{a}\eta = a$. We will usually write \bar{a} simply as a . An example of constants that will be useful later are $\bar{\mathbf{0}}$ and $\bar{\mathbf{1}}$ that represent respectively the constraints associating $\mathbf{0}$ and $\mathbf{1}$ to all the assignment of domain values.

Solutions. A SCSP (Bistarelli 2004) is defined as $P = \langle V, D, C, S \rangle$, where C is the set of constraints defined over variables in V (each with domain D), and whose preference is determined by semiring S . The *best level of consistency* notion is defined as $blevel(P) = Sol(P) \Downarrow_{\emptyset}$, where $Sol(P) = \bigotimes C$ (Bistarelli 2004). A problem P is α -consistent if $blevel(P) = \alpha$ (Bistarelli 2004). P is instead simply “consistent” iff there exists $\alpha >_S \mathbf{0}$ such that P is α -consistent. P is inconsistent if it is not consistent.




Fig. 1. A SCSP based on a weighted semiring.

Example 1

Figure 1 shows a weighted SCSP as a graph: the weighted semiring is used, i.e. $\langle \mathbb{R}^+ \cup \{\infty\}, \min, \hat{+}, \infty, 0 \rangle$ ($\hat{+}$ is the arithmetic plus operation). Variables and constraints are represented respectively by nodes and arcs (unary for c_1 - c_3 , and binary for c_2); $D = \{a, b\}$. The solution of the CSP in Figure 1 associates a semiring element to

every domain value of variables X and Y by combining all the constraints together, i.e. $Sol(P) = \bigotimes C$. For instance, for the tuple $\langle a, a \rangle$ (that is, $X = Y = a$), we have to compute the sum of 1 (which is the value assigned to $X = a$ in constraint c_1), 5 ($\langle X = a, Y = a \rangle$ in c_2) and 5 ($Y = a$ in c_3): the value for this tuple is 11. The solution $X = a, Y = b$ is a 7-consistent solution, where 7 corresponds to the *blevel* of P , i.e., $Sol(P) \Downarrow_{\emptyset} = 7$.

2.2 Concurrent Constraint Programming over Soft Constraints

The basic idea underlying *ccp* (Saraswat 1989) is that computation progresses via monotonic accumulation of information in a constraint global store. Information is produced by the concurrent and asynchronous activity of several agents which can add (*tell*) a constraint to the store. Dually, agents can also check (*ask*) whether a constraint is entailed by the store, thus allowing synchronization among different agents. The *ccp* languages are defined parametrically w.r.t. a given *constraint system*. The notion of constraint system has been formalized in (Saraswat and Rinard 1990) following Scott's treatment of information systems. Soft constraints over a semiring $S = \langle \mathcal{A}, +, \times, \mathbf{0}, \mathbf{1} \rangle$ and an ordered set of variables V (over a domain D) have been showed to form a constraint system “à la Saraswat”, thus leading to the definition of *Soft Concurrent Constraint Programming* (*sccp*) (Bistarelli et al. 1997; Bistarelli 2004; Bistarelli et al. 2006).

Consider the set \mathcal{C} and the partial order \sqsubseteq . Then an entailment relation $\vdash \subseteq \wp(\mathcal{C}) \times \mathcal{C}$ is defined s.t. for each $C \in \wp(\mathcal{C})$ and $c \in \mathcal{C}$, we have $C \vdash c \Leftrightarrow \bigotimes C \sqsubseteq c$ (see also (Bistarelli 2004; Bistarelli et al. 2006)). Note that in this setting the notion of token (constraint) and of set of tokens (set of constraints) closed under entailment is used indifferently. In fact, given a set of constraint functions C_1 , its closure w.r.t. entailment is a set \bar{C}_1 that contains all the constraints greater than $\bigotimes C_1$. This set is univocally representable by the constraint function $\bigotimes C_1$. The definition of the entailment operator \vdash on top of \mathcal{C} , and of the \sqsubseteq relation, lead to the notion of *soft constraint system*. It is also important to notice that in (Saraswat 1989) it is claimed that a constraint system is a *complete algebraic* lattice. In the *sccp* framework, algebraicity is not required (Bistarelli et al. 2006) instead, since the algebraic nature of the structure \mathcal{C} strictly depends on the properties of the semiring¹.

To treat the hiding operator of the language, a general notion of existential quantifier is introduced by using notions similar to those used in cylindric algebras. Consider a set of variables V with domain D and the corresponding soft constraint system \mathcal{C} . For each $x \in V$, the hiding function (Bistarelli 2004; Bistarelli et al. 2006) is the function $(\exists_x c)\eta = \sum_{d_i \in D} c\eta[x := d_i]$. To make the hiding operator computationally tractable, it is required that the number of domain elements in D , having semiring values different from $\mathbf{0}$, is finite (Bistarelli et al. 2006). In this way, to compute the sum needed for $(\exists_x c)\eta$, we can consider just a finite number of elements

¹ Notice that we do not aim at computing the closure of the entailment relation, but only to use the entailment relation to establish if a constraint is entailed by the current store, and this can be established even if the lattice is not algebraic (that is even if the times operator is not idempotent).

(those different from $\mathbf{0}$), since $\mathbf{0}$ is the unit element of the sum. Note that by using the hiding function we can represent the \Downarrow operator defined in Section 2.1. In fact, for any constraint c and any variable $x \subseteq V$, $c \Downarrow_{V-x} = \exists_x c$ (Bistarelli et al. 2006).

To model parameter passing also diagonal elements have to be defined. Consider a set of variables V and the corresponding soft constraint system. Then, for each $x, y \in V$, a diagonal constraint is defined as $d_{xy} \in \mathcal{C}$ s.t., $d_{xy}\eta[x := a, y := b] = \mathbf{1}$ if $a = b$, and $d_{xy}\eta[x := a, y := b] = \mathbf{0}$ if $a \neq b$ (Bistarelli 2004; Bistarelli et al. 2006).

Theorem 1 (cylindric constraint system (Bistarelli et al. 2006))

Consider a semiring $S = \langle \mathcal{A}, +, \times, \mathbf{0}, \mathbf{1} \rangle$, a domain of the variables D , an ordered set of variables V , and the corresponding structure \mathcal{C} . Then, $S_C = \langle \mathcal{C}, \otimes, \mathbf{0}, \mathbf{1}, \exists_x, d_{xy} \rangle$, is a cylindric constraint system.

2.3 Timed Concurrent Constraint Programming

A timed extension of *ccp*, called *tccp* has been introduced in (de Boer et al. 2000). Similarly to other existing timed extensions of *ccp* defined in (Saraswat et al. 1996), *tccp* is a language for reactive programming designed around the hypothesis of *bounded asynchrony* (as introduced in (Saraswat et al. 1996): computation takes a bounded period of time rather than being instantaneous).

When querying the store for some information that is not present (yet), a *ccp* agent will simply suspend until the required information has arrived. In timed applications however often one cannot wait indefinitely for an event. Consider for example the case of a connection to a web service providing some on-line banking facility. In case the connection cannot be established, after a reasonable amount of time an appropriate time-out message has to be communicated to the user. A timed language should then allow us to specify that, in case a given time bound is exceeded (i.e. a *time-out* occurs), the wait is interrupted and an alternative action is taken. Moreover, in some cases it is also necessary to have a preemption mechanism which allows one to abort an active process A and to start a process B when a specific (abnormal) event occurs.

In order to be able to specify these timing constraints *tccp* introduces a discrete global clock and assumes that *ask* and *tell* actions take one time-unit. Computation evolves in steps of one time-unit, so called clock-cycles. Action prefixing is the syntactic marker which distinguishes a time instant from the next one and it is assumed that parallel processes are executed on different processors, which implies that, at each moment, every enabled agent of the system is activated. This assumption gives rise to what is called *maximal parallelism*. The time in between two successive moments of the global clock intuitively corresponds to the response time of the underlying constraint system. Thus all parallel agents are synchronized by the response time of the underlying constraint system. Since the store is monotonically increasing and one can have dynamic process creation, clearly the previous assumptions imply that the constraint solver takes a constant time (no matter how big the store is), and that there is an unbounded number of processors. However, one can impose suitable restriction on programs, thus ensuring that the (significant

part of the) store and the number of processes do not exceed a fixed bound; these restrictions would still allow significant forms of recursion with parameters.

Furthermore, a timing construct of the form **now** c **then** A **else** B is introduced in *tccp*, whose semantics is the following: if the constraint c is entailed by the store at the current time t , then the above agent behaves as A at time t , otherwise it behaves as B at time t . This basic construct allows to derive such timing mechanisms as time-out and preemption (de Boer et al. 2000; Saraswat et al. 1996). The instantaneous reaction can be obtained by evaluating **now** c in parallel with A and B , within the same time-unit. At the end of this time-unit, the store will be updated by using either the constraint produced by A , or that one produced by B , depending on the result of the evaluation of **now** c . Clearly, since A and B could contain nested **now then else** agents, a limit for the number of these nested agents should be fixed. Note that, for recursive programs, such a limit is ensured by the presence of the procedure-call, since we assume that the evaluation of such calls takes one time-unit.

3 Timed Soft Concurrent Constraint Programming

In this section we present the *tsccp* language, which originates from both *tccp* and *sccp*. To obtain this aim, we extend the syntax of the *cc* language with the timing construct **now** c **then** A **else** B (inherited from *tccp*), and also in order to directly handle the cut level as in *sccp*. This means that the syntax and semantics of the **tell**, **ask** and **now** agents have to be enriched with a threshold that is used to check when the agents may succeed, or suspend.

Definition 1 (tsccp Language)

Given a soft constraint system $\langle S, D, V \rangle$, the corresponding structure \mathcal{C} , any semiring value a , soft constraints ϕ , $c \in \mathcal{C}$ and any tuple of variables x , the syntax of the *tsccp* language is given by the following grammar:

$$\begin{aligned}
P &::= F.A \\
F &::= p(x) :: A \mid F \cdot F \\
A &::= \mathbf{success} \mid \mathbf{tell}(c) \rightarrow_{\phi} A \mid \mathbf{tell}(c) \rightarrow^a A \mid E \mid A \parallel A \mid \exists x A \mid p(x) \mid \\
&\quad \Sigma_{i=1}^n E_i \mid \mathbf{now}_{\phi} c \mathbf{then} A \mathbf{else} A \mid \mathbf{now}^a c \mathbf{then} A \mathbf{else} A \\
E &::= \mathbf{ask}(c) \rightarrow_{\phi} A \mid \mathbf{ask}(c) \rightarrow^a A
\end{aligned}$$

where, as usual, P is the class of processes, F is the class of sequences of procedure declarations (or clauses), A is the class of agents. In a *tsccp* process $P = F.A$, A is the initial agent, to be executed in the context of the set of declarations F . The agent **success** represents a successful termination, so it may not make any further transition.

In the following, given an agent A , we denote by $Fv(A)$ the set of the free variables of A (namely, the variables which do not appear in the scope of the \exists quantifier). Besides the use of soft constraints (see Section 2.2) instead of crisp ones, there are two fundamental differences between *tsccp* and *ccp*. The first main difference w.r.t. the original *cc* syntax is the presence of a semiring element a and of a constraint

ϕ to be checked whenever an *ask* or *tell* operation is performed. More precisely, the level a (respectively, ϕ) will be used as a cut level to prune computations that are not good enough. The second main difference with respect to *ccp* (but, this time, also with respect to *sccp*) is instead the presence of the **nowc then A else B** construct introduced in Section 2.3. Even for this construct, the level a (or ϕ) is used as a cut level to prune computations.

Action prefixing is denoted by \rightarrow , non-determinism is introduced via the guarded choice construct $\Sigma_{i=1}^n E_i$, parallel composition is denoted by \parallel , and a notion of locality is introduced by the agent $\exists xA$, which behaves like A with x considered local to A , thus hiding the information on x provided by the external environment.

In the next subsection we formally describe the operational semantics of *tsccp*. In order to simplify the notation, in the following we will usually write a *tsccp process* $P = F.A$ simply as the corresponding agent A .

4 An Operational Semantics for *tsccp* Agents

The operational model of *tsc* agents can be formally described by a transition system $T = (Conf, \rightarrow)$ where we assume that each transition step takes exactly one time-unit. Configurations in *Conf* are pairs consisting of a process and of a constraint in \mathcal{C} , representing the common *store* shared by all the agents. The transition relation $\rightarrow \subseteq Conf \times Conf$ is the least relation satisfying the rules **R1-R17** in Figure 2, and it characterizes the (temporal) evolution of the system. So, $\langle A, \sigma \rangle \rightarrow \langle B, \delta \rangle$ means that, if at time t we have the process A and the store σ , then at time $t + 1$ we have the process B and the store δ .

Let us now briefly discuss the rules in Figure 2. Here is a brief description of the transition rules:

Valued-tell. The valued-tell rule checks for the a -consistency of the *Soft Constraint Satisfaction Problem* (Bistarelli 2004) (SCSP) defined by the store $\sigma \otimes c$. A SCSP P is a -consistent if $blevel(P) = a$, where $blevel(P) = Sol(P) \Downarrow_{\emptyset}$, i.e., the *best level of consistency* of the problem P is a semiring value representing the least upper bound among the values yielded by the solutions. Rule **R1** can be applied only if the store $\sigma \otimes c$ is b -consistent with $b \not\leq a^2$. In this case the agent evolves to the new agent A over the store $\sigma \otimes c$. Note that different choices of the *cut level* a could possibly lead to different computations. Finally, note that the updated store $\sigma \otimes c$ will be visible only starting from the next time instant, since each transition step involves exactly one time-unit.

Tell. The tell action is a finer check of the store. In this case (see rule **R2**), a pointwise comparison between the store $\sigma \otimes c$ and the constraint ϕ is performed. The idea is to perform an overall check of the store, and to continue the computation only if there is the possibility to compute a solution not worse than ϕ . Note that this notion of tell could be also applied to the classical *cc* framework:

² Notice that we use $b \not\leq a$ instead of $b \geq a$ because we can possibly deal with partial orders. The same holds also for $\not\sqsubseteq$ instead of \sqsupseteq .

R1	$\frac{(\sigma \otimes c) \Downarrow_{\emptyset} \not\prec a}{\langle \text{tell}(c) \rightarrow^a A, \sigma \rangle \longrightarrow \langle A, \sigma \otimes c \rangle}$	V-tell
R2	$\frac{\sigma \otimes c \not\sqsubseteq \phi}{\langle \text{tell}(c) \rightarrow_{\phi} A, \sigma \rangle \longrightarrow \langle A, \sigma \otimes c \rangle}$	Tell
R3	$\frac{\sigma \vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a}{\langle \text{ask}(c) \rightarrow^a A, \sigma \rangle \longrightarrow \langle A, \sigma \rangle}$	V-ask
R4	$\frac{\sigma \vdash c \quad \sigma \not\sqsubseteq \phi}{\langle \text{ask}(c) \rightarrow_{\phi} A, \sigma \rangle \longrightarrow \langle A, \sigma \rangle}$	Ask
R5	$\frac{\langle A, \sigma \rangle \longrightarrow \langle A', \sigma \otimes \delta \rangle \quad \langle B, \sigma \rangle \longrightarrow \langle B', \sigma \otimes \delta' \rangle}{\langle A \parallel B, \sigma \rangle \longrightarrow \langle A' \parallel B', \sigma \otimes \delta \otimes \delta' \rangle}$	Parall1
R6	$\frac{\langle A, \sigma \rangle \longrightarrow \langle A', \sigma' \rangle \quad \langle B, \sigma \rangle \not\rightarrow}{\langle A \parallel B, \sigma \rangle \longrightarrow \langle A' \parallel B, \sigma' \rangle}$ $\langle B \parallel A, \sigma \rangle \longrightarrow \langle B \parallel A', \sigma' \rangle$	Parall2
R7	$\frac{\langle E_j, \sigma \rangle \longrightarrow \langle A_j, \sigma' \rangle \quad j \in [1, n]}{\langle \sum_{i=1}^n E_i, \sigma \rangle \longrightarrow \langle A_j, \sigma' \rangle}$	Nondet
R8	$\frac{\langle A, \sigma \rangle \longrightarrow \langle A', \sigma' \rangle \quad \sigma \vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a}{\langle \text{now}^a c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle A', \sigma' \rangle}$	V-now1
R9	$\frac{\langle A, \sigma \rangle \not\rightarrow \quad \sigma \vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a}{\langle \text{now}^a c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle A, \sigma \rangle}$	V-now2
R10	$\frac{\langle B, \sigma \rangle \longrightarrow \langle B', \sigma' \rangle \quad \sigma \not\vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a}{\langle \text{now}^a c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle B', \sigma' \rangle}$	V-now3
R11	$\frac{\langle B, \sigma \rangle \not\rightarrow \quad \sigma \not\vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a}{\langle \text{now}^a c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle B, \sigma \rangle}$	V-now4
R12	$\frac{\langle A, \sigma \rangle \longrightarrow \langle A', \sigma' \rangle \quad \sigma \vdash c \quad \sigma \not\sqsubseteq \phi}{\langle \text{now}_{\phi} c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle A', \sigma' \rangle}$	Now1
R13	$\frac{\langle A, \sigma \rangle \not\rightarrow \quad \sigma \vdash c \quad \sigma \not\sqsubseteq \phi}{\langle \text{now}_{\phi} c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle A, \sigma \rangle}$	Now2
R14	$\frac{\langle B, \sigma \rangle \longrightarrow \langle B', \sigma' \rangle \quad \sigma \not\vdash c \quad \sigma \not\sqsubseteq \phi}{\langle \text{now}_{\phi} c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle B', \sigma' \rangle}$	Now3
R15	$\frac{\langle B, \sigma \rangle \not\rightarrow \quad \sigma \not\vdash c \quad \sigma \not\sqsubseteq \phi}{\langle \text{now}_{\phi} c \text{ then } A \text{ else } B, \sigma \rangle \longrightarrow \langle B, \sigma \rangle}$	Now4
R16	$\frac{\langle A[x/y], \sigma \rangle \longrightarrow \langle B, \sigma' \rangle}{\langle \exists x A, \sigma \rangle \longrightarrow \langle B, \sigma' \rangle}$	Hide
R17	$\langle p(x), \sigma \rangle \longrightarrow \langle A, \sigma \rangle \quad p(x) :: A \in F$	P-call

Fig. 2. The transition system for *tsccp*.

the tell operation would succeed when the set of tuples satisfying constraint ϕ is not a superset of the set of tuples allowed by $\sigma \cap c$.³ As for the valued tell, the updated store $\sigma \otimes c$ will be visible only since the next time instant. In the following, let us use $\mathbf{tell}(c) \rightarrow A$ and $\mathbf{tell}(c)$ as a shorthand for $\mathbf{tell}(c) \rightarrow_{\mathbf{0}} A$ and $\mathbf{tell}(c) \rightarrow_{\mathbf{0}} \mathbf{success}$, respectively.

Valued-ask. The semantics of the valued-ask is extended in a way similar to what we have done for the valued-tell action. This means that, to apply the rule **R3**, we need to check if the store σ entails the constraint c , and also if σ is “consistent enough” w.r.t. the threshold a set by the programmer.

Ask. In rule **R4**, we check if the store σ entails the constraint c , but, similarly to rule **R2**, we also compare a finer (pointwise) threshold ϕ to the store σ . As for the tell action, let us use $\mathbf{ask}(c) \rightarrow A$ as a shorthand for $\mathbf{ask}(c) \rightarrow_{\mathbf{0}} A$.

Parallelism. Rules **R5** and **R6** model the parallel composition operator in terms of *maximal parallelism*: the agent $A \parallel B$ executes in one time-unit all the initial enabled actions of A and B . Considering rule **R5** (where *maximal parallelism* is accomplished in practice), notice that the ordering of the operands in $\sigma \otimes \delta \otimes \delta'$ is not relevant, since \otimes is commutative and associative. Moreover, for the same two properties, if $\sigma \otimes \delta = \sigma \otimes \gamma$ and $\sigma \otimes \delta' = \sigma \otimes \gamma'$, we have that $\sigma \otimes \delta \otimes \delta' = \sigma \otimes \gamma \otimes \gamma'$. Therefore the resulting store $\sigma \otimes \delta \otimes \delta'$ is independent from the choice of the constraint δ such that $\langle A, \sigma \rangle \longrightarrow \langle A', \sigma' \rangle$ and $\sigma' = \sigma \otimes \delta$ (analogously for δ').

Nondeterminism. According to rule **R7**, the guarded choice operator gives rise to global non-determinism: the external environment can affect the choice, since $\mathbf{ask}(c_j)$ is enabled at time t (and A_j is started at time $t + 1$) if and only if the store σ entails c_j (and if it is compatible with the threshold too), and σ can be modified by other agents.

Valued-now and Now. Rules **R8-R11** show that the agent $\mathbf{now}^a c \mathbf{then} A \mathbf{else} B$ behaves as A or B depending on the fact that c is or is not entailed by the store, provided that the current store σ is compatible with the threshold. Differently from the case of the ask, here the evaluation of the guard is instantaneous: if current store σ is compatible with the threshold a , $\langle A, \sigma \rangle$ ($\langle B, \sigma \rangle$) can make a transition at time t and c is (is not) entailed by the store σ , then the agent $\mathbf{now}^a c \mathbf{then} A \mathbf{else} B$ can make the same transition at time t . Moreover, observe that in any case the control is passed either to A (if c is entailed by the current store σ and σ is compatible with the threshold) or to B (in case σ does not entail c and σ is compatible with the threshold). Analogously for the not-valued version, i.e., $\mathbf{now}_\phi c \mathbf{then} A \mathbf{else} B$ (see rules **R12-R15**). Finally, we use $\mathbf{now} c \mathbf{then} A \mathbf{else} B$ as a shorthand for the agent $\mathbf{now}_{\mathbf{0}} c \mathbf{then} A \mathbf{else} B$.

Hiding variables. The agent $\exists x A$ behaves like A , with x considered *local* to A , as show by rule **R16**. This is obtained by substituting the variable x for a variable y , which we assume to be new and not used by any other process. Standard renaming techniques can be used to ensure this; in rule **R16**, $A[x/y]$ denotes the process obtained from A by replacing the variable x for the variable y .

³ Notice that the \otimes operator in the crisp case reduces to set intersection.

Procedure-calls. Rule **R17** treats the case of a procedure-call when the actual parameter equals the formal parameter. We do not need more rules since, for the sake of simplicity, here and in the following we assume that the set F of procedure declarations is closed w.r.t. parameter names: that is, for every procedure-call $p(y)$ appearing in a process $F.A$, we assume that, if the original declaration for p in F is $p(x) :: A$, then F contains also the declaration $p(y) :: \exists x(\mathbf{tell}(d_{xy}) \parallel A)$.⁴ Moreover, we assume that if $p(x) :: A \in F$, then $Fv(A) \subseteq x$.

Using the transition system described by (the rules in) Figure 2, we can now define our notion of observables, which considers the results of successful terminating computations that the agent A can perform for each *tsccp* process $P = F.A$. Here and in the following, given a transition relation \longrightarrow , we denote by \longrightarrow^* its reflexive and transitive closure.

Definition 2 (Observables)

Let $P = F.A$ be a *tsccp* process. We define

$$\mathcal{O}_{io}^{mp}(P) = \{\gamma \Downarrow_{Fv(A)} \mid \langle A, \mathbf{1} \rangle \longrightarrow^* \langle \mathbf{Success}, \gamma \rangle\}$$

where **Success** is any agent which contains only occurrences of the agent **success** and of the operator \parallel .

5 Programming Idioms and Examples

We can consider the primitives in Definition 1 to derive the soft version of the programming idioms in (de Boer et al. 2000), which are typical of reactive programming.

Delay. The delay constructs $\mathbf{tell}(c) \xrightarrow{t}_\phi A$ or $\mathbf{ask}(c) \xrightarrow{t}_\phi A$ are used to delay the execution of agent A after the execution of $\mathbf{tell}(c)$ or $\mathbf{ask}(c)$; t is the number of the time-units of delay. Therefore, in addition to a constraint ϕ , in *tsccp* the transition arrow can have also a number of delay slots. This idiom can be defined by induction: the base case is $\xrightarrow{0}_\phi A \equiv \rightarrow_\phi A$, and the inductive step is $\xrightarrow{n+1}_\phi A \equiv \rightarrow_\phi \mathbf{tell}(\bar{1}) \xrightarrow{n}_{\bar{0}} A$. The valued version can be defined in an analogous way.

Timeout. The timed guarded choice agent $\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i \mathbf{timeout}(m) B$ waits at most m time-units ($m \geq 0$) for the satisfaction of one of the guards; notice that all the ask actions have a soft transition arrow, i.e. \rightarrow_i is either of the form \rightarrow_{ϕ_i} or \rightarrow^{a_i} , as in Figure 2. Before this time-out, the process behaves just like the guarded choice: as soon as there exist enabled guards, one of them (and the corresponding branch) is nondeterministically selected. After waiting for m time-units, if no guard is enabled, the timed choice agent behaves as B . Timeout constructs can be assembled through the composition of several

⁴ Here the (original) formal parameter is identified as a local alias of the actual parameter. Alternatively, we could have introduced a new rule treating explicitly this case, as it was in the original *ccp* papers.

now_φ *c* **then** *A* **else** *B* primitives (or their valued version), as explained in (de Boer et al. 2000) for the (crisp) *tccp* language.

The timeout can be defined inductively as follows: let us denote by *A* the agent $\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i$. In the base case, that is $m = 0$, we define $\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i \mathbf{timeout}(0) B$ as the agent:

$$\begin{aligned} & \mathbf{now}_1 c_1 \mathbf{then} A \\ & \quad \mathbf{else} (\mathbf{now}_2 c_2 \mathbf{then} A \\ & \quad \quad \mathbf{else} (\dots (\mathbf{now}_n c_n \mathbf{then} A \mathbf{else} \mathbf{ask}(\bar{1}) \rightarrow B) \dots)) \end{aligned}$$

where for $i = 1, \dots, n$, either $\mathbf{now}_i = \mathbf{now}_{\phi_i}$ if \rightarrow_i is of the form \rightarrow_{ϕ_i} or $\mathbf{now}_i = \mathbf{now}^{a_i}$ if \rightarrow_i is of the form \rightarrow^{a_i} . Because of the operational semantics explained in rules **R8-R11** (see Figure 2), if a guard c_i is true, then the agent $\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i$ is evaluated in the same time slot. Otherwise, if no guard c_i is true, the agent *B* is evaluated in the next time slot. Then, by inductively reasoning on the number of time-units m , we can define $\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i \mathbf{timeout}(m) B$ as

$$\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i \mathbf{timeout}(0) (\Sigma_{i=1}^n \mathbf{ask}(c_i) \rightarrow_i A_i \mathbf{timeout}(m-1) B).$$

Watchdog. Watchdogs are used to interrupt the activity of a process on a signal from a specific event. The idiom **do** *A* **watching**_φ *c* behaves as *A*, as long as *c* is not entailed by the store and the current store is compatible with the threshold; when *c* is entailed and the current store is compatible with the threshold, the process *A* is immediately aborted.

The reaction is instantaneous, in the sense that *A* is aborted at the same time instant of the detection of the entailment of *c*. However, according to the computational model, if *c* is detected at time t , then *c* has to be produced at time t' with $t' < t$. Thus, we have a form of weak preemption.

As well as timeouts, also watchdog agents can be defined in terms of the other basic constructs of the language (see Figure 3).

In the following we assume that there exists an (injective) renaming function ρ which, given a procedure name p , returns a new name $\rho(p)$ that is not used elsewhere in the program. Moreover, let us use **now**_φ *c* **else** *B* as a shorthand for **now**_φ *c* **then success** **else** *B*, where we assume that, for any procedure p declared as $p(x) :: A$, a declaration $\rho(p)(x) :: \mathbf{do} \rho(A) \mathbf{watching}_{\phi} c$ is added, where $\rho(A)$ denotes the agent obtained from *A* by replacing in it each occurrence of any procedure q by $\rho(q)$. The assumption in the case of the $\exists x A$ agent is needed for correctness. In practical cases, it can be satisfied by suitably renaming the variables associated to signals. In the following \rightarrow' is either of the form \rightarrow_{ψ} or \rightarrow^a . Analogously for **now**'.

The translation in Figure 3 can be easily extended to the case of the agent **do** *A* **watching**_φ *c* **else** *B*, which behaves as the previous watchdog and also activates the process *B* when *A* is aborted (i.e., when *c* is entailed and the current state is compatible with the threshold). In the following we will then use also this form of watchdog.

do success $\text{watching}_\phi c$	\implies	success
do tell (d) $\rightarrow' A$ $\text{watching}_\phi c$	\implies	now $_\phi c$ else tell (d) $\rightarrow' \text{do } A$ $\text{watching}_\phi c$
do $\Sigma_{i=1}^n \text{ask}(c_i) \rightarrow_i A_i$ $\text{watching}_\phi c$	\implies	now $_\phi c$ else $\Sigma_{i=1}^n \text{ask}(c_i) \rightarrow_i \text{do } A_i$ $\text{watching}_\phi c$
do (now' d then A else B) $\text{watching}_\phi c$	\implies	now' d then do A $\text{watching}_\phi c$ else do B $\text{watching}_\phi c$
do $A \parallel B$ $\text{watching}_\phi c$	\implies	do A $\text{watching}_\phi c \parallel \text{do } B$ $\text{watching}_\phi c$
do $\exists x A$ $\text{watching}_\phi c$	\implies	$\exists x$ do A $\text{watching}_\phi c$, assuming $\exists_x c = c$
do $p(x)$ $\text{watching}_\phi c$	\implies	now $_\phi c$ else $\rho(p)(x)$ $\text{watching}_\phi c$

Fig. 3. Examples of watchdog constructs.

$$\begin{aligned}
c_1 : (\{x\} \rightarrow N) \rightarrow R^+ \quad \text{s.t. } c_1(x) = x + 3 & \quad c_2 : (\{x\} \rightarrow N) \rightarrow R^+ \quad \text{s.t. } c_2(x) = x + 5 \\
c_3 : (\{x\} \rightarrow N) \rightarrow R^+ \quad \text{s.t. } c_3(x) = 2x + 8
\end{aligned}$$

Fig. 4. Three (weighted) soft constraints; $c_3 = c_1 \otimes c_2$, $c_2 \vdash c_1$, $c_3 \vdash c_1$ and $c_3 \vdash c_2$.

The assumption on the instantaneous evaluation of **now** $_\phi c$ is essential in order to obtain a preemption mechanism which can be expressed in terms of the **now** $_\phi$ **then else** primitive. In fact, if the evaluation of **now** $_\phi c$ took one time-unit, then this unit delay would change the compositional behavior of the agent controlled by the watchdog. Consider, for example, the agent $A = \text{tell}(a) \rightarrow \text{tell}(b)$, which takes two time-units to complete its computation. The agent $A^t = \text{now } c \text{ else tell}(a) \rightarrow \text{now } c \text{ else tell}(b)$ (resulting from the translation of **do** A **watching** $_{\bar{0}} c$) compositionally behaves as A , unless a c signal is detected and the current state is compatible with the threshold, in which case the evaluation of A is interrupted. On the other hand, if the evaluation of **now** c took one time-unit, then A^t would take four time-units and would not behave anymore as A when c is not present. In fact, in this case, the agent $A \parallel B$ would produce d while $A^t \parallel B$ would not, where B is the agent $\text{ask}(\bar{1}) \rightarrow \text{now } a \text{ then tell}(d) \text{ else success}$.

The valued version of watchdogs can be defined in an analogous way.

With this small set of idioms, we have now enough expressiveness to describe complex interactions. For the following examples on the new programming idioms, we consider the *Weighted* semiring $\langle \mathbb{R}^+ \cup \{+\infty\}, \min, +, +\infty, 0 \rangle$ (Bistarelli 2004; Bistarelli et al. 1997) and the (weighted) soft constraints in Figure 4. We first provide simple program examples in order to explain as more details as possible on how a computation of *tsccp* agents proceeds. In Section 5.1 we show a more complex example describing the classical actions during a negotiation process; the aim of that example is instead to show the expressivity of the *tsccp* language, without analyzing its execution in detail.

Example 2 (Delay)

As a first very simple example, suppose to have two agents A_1, A_2 of the form: $A_1 :: \mathbf{tell}(\bar{1}) \xrightarrow{2} +\infty \mathbf{tell}(c_2) \rightarrow +\infty \mathbf{success}$ and $A_2 :: \mathbf{tell}(\bar{1}) \xrightarrow{1} +\infty \mathbf{ask}(c_1) \rightarrow^9 \mathbf{success}$; their concurrent evaluation in the $\bar{1} \equiv \bar{0}$ empty store is:

$$\langle (\mathbf{tell}(\bar{0}) \xrightarrow{2} +\infty \mathbf{tell}(c_2) \rightarrow +\infty \mathbf{success}) \parallel (\mathbf{tell}(\bar{0}) \xrightarrow{1} +\infty \mathbf{ask}(c_1) \rightarrow^9 \mathbf{success}), \bar{0} \rangle.$$

The timeline for this parallel execution is described in Figure 5. For the evaluation of **tell** and **ask** we respectively consider the rules **R1** and **R3** in Figure 2, since both transitions are a -valued. However, both these two actions are delayed: three time-units for the **tell**(c_2) of A_1 (including the first **tell**($\bar{0}$)), and two time-units for the **ask**(c_1) of A_2 (including the first **tell**($\bar{0}$)). As explained before, this can be obtained by adding $\bar{1}$ to the store with a **tell** action respectively three, and two times. Therefore, the parallel agent $A_1 \parallel A_2$ corresponds to:

$$(\mathbf{tell}(\bar{0}) \rightarrow +\infty \mathbf{tell}(\bar{0}) \rightarrow +\infty \mathbf{tell}(\bar{0}) \rightarrow +\infty \mathbf{tell}(c_2) \rightarrow +\infty \mathbf{success}) \parallel (\mathbf{tell}(\bar{0}) \rightarrow +\infty \mathbf{tell}(\bar{0}) \rightarrow +\infty \mathbf{ask}(c_1) \rightarrow^9 \mathbf{success}).$$

This agent is interpreted by using **R5-R6** in Figure 2 in terms of maximal parallelism, i.e., all the actions are executed in parallel. The first two **tell** of A_1 and A_2 can be simultaneously executed by using rule **R1**: the precondition $(\bar{0} \otimes \bar{0}) \Downarrow_0 = 0 \not\prec 9$ of the rule is then satisfied. The store does not change since $\bar{0} \otimes \bar{0} = \bar{0}$. At this point, the **ask** action of A_2 is not enabled because $\bar{0} \not\vdash c_1$, that is the precondition $\sigma \vdash c_1$ of **R3** is not satisfied. Therefore, the processor can only be allocated to A_1 and, since $(\bar{0} \otimes \bar{0}) \Downarrow_0 = 0 \not\prec +\infty$ is true (i.e. the precondition of **R1** is satisfied), at $t = 3$ the computation is in the state:

$$\langle \mathbf{tell}(c_2) \rightarrow +\infty \mathbf{success} \parallel \mathbf{ask}(c_1) \rightarrow^9 \mathbf{success}, \bar{0} \rangle.$$

Now the **tell** can be executed because $(\bar{0} \otimes c_2) \Downarrow_0 = 5 \not\prec +\infty$: therefore, the store becomes equal to $\bar{0} \otimes c_2 = c_2$:

$$\langle \mathbf{success} \parallel \mathbf{ask}(c_1) \rightarrow^9 \mathbf{success}, c_2 \rangle.$$

At $t = 5$ (see Figure 5) we can successfully terminate the program: in the store $\sigma = c_2$ the **ask** is finally enabled at $t = 4$, according to the two preconditions of rule **R3**, i.e., $c_2 \vdash c_1$ and $c_2 \Downarrow_0 = 5 \not\prec 9$: therefore we have $A_1 \parallel A_2 :: \langle \mathbf{success} \parallel \mathbf{success}, c_2 \rangle$.

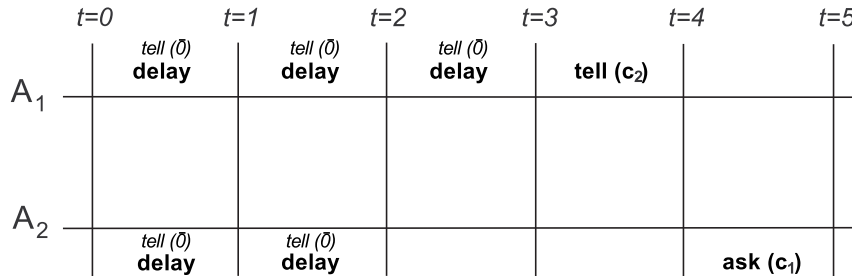


Fig. 5. The timeline of the execution of the $A_1 \parallel A_2$ parallel agent in Example 2.

Example 3 (Timeout)

In this second example we evaluate a timeout construct. Suppose we have two agents A_1 and A_2 of the form:

$$A_1 :: ((\mathbf{ask}(c_1) \rightarrow^{+\infty} \mathbf{success}) + (\mathbf{ask}(c_2) \rightarrow^{+\infty} \mathbf{success})) \mathbf{timeout}(1)$$

$$\mathbf{ask}(c_1) \rightarrow^{+\infty} \mathbf{success}$$

and

$$A_2 :: \mathbf{tell}(\bar{0}) \xrightarrow{2}^{+\infty} \mathbf{tell}(c_3) \rightarrow^{+\infty} \mathbf{success}$$

The description of agent A_1 is a shortcut for the following agent, as previously explained in the definition of the timeout:

$$\mathbf{now}^{+\infty} c_1 \mathbf{then} B \mathbf{else} (\mathbf{now}^{+\infty} c_2 \mathbf{then} B \mathbf{else}$$

$$(\mathbf{ask}(\bar{1}) \rightarrow \mathbf{now}^{+\infty} c_1 \mathbf{then} B \mathbf{else}$$

$$(\mathbf{now}^{+\infty} c_2 \mathbf{then} B \mathbf{else} (\mathbf{ask}(\bar{1}) \rightarrow \mathbf{ask}(c_1) \rightarrow^{+\infty} \mathbf{success}))))).$$

where $B :: (\mathbf{ask}(c_1) \rightarrow^{+\infty} \mathbf{success} + \mathbf{ask}(c_2) \rightarrow^{+\infty} \mathbf{success})$. Their concurrent evaluation in the $\bar{1} \equiv \bar{0}$ empty store is:

$$\langle (B \mathbf{timeout}(1) \mathbf{ask}(c_1) \rightarrow^{+\infty} \mathbf{success} \parallel \mathbf{tell}(\bar{0}) \xrightarrow{2}^{+\infty} \mathbf{tell}(c_3) \rightarrow^{+\infty} \mathbf{success}), \bar{0} \rangle.$$

The timeline for this parallel execution is given in Figure 6. At $t = 0$ the store is empty (i.e., $\sigma = \bar{0}$), thus both constraints c_1 and c_2 asked by the nondeterministic choice agent A_1 are not entailed. In A_2 , the \mathbf{tell} of c_3 , which would entail both c_1 and c_2 , is delayed by three time-units: in the first three time-units, $\mathbf{tell}(\bar{0}) \rightarrow^{+\infty}$ is executed according to the delay construct, as shown in Example 2. At $t = 2$ the timeout is triggered in A_1 , since, according to **R1**, **R6** and **R9** (see Figure 2), the time elapsing in the timeout construct can be executed together with the delay- \mathbf{tell} actions of A_2 . After the timeout triggering, agent A_1 is however blocked, since c_1 is not entailed by the current empty store, and the precondition of the \mathbf{ask} (rule **R3**) is not satisfied. A_2 can execute the last delay- \mathbf{tell} , and then perform the $\mathbf{tell}(c_3)$ operation at $t = 3$; the store becomes $\sigma = \bar{0} \otimes c_3 = c_3$. This finally unblocks A_1 at $t = 4$, since, according to the precondition of rule **R3**, $\sigma \sqsubseteq c_1$ (i.e., $c_3 \sqsubseteq c_1$). Finally, at $t = 5$ we have $\langle \mathbf{success} \parallel \mathbf{success}, c_3 \rangle$.




Fig. 6. The timeline of the execution of the $A_1 \parallel A_2$ agent in Example 3.




Fig. 7. The timeline of the execution of the $A_1 \parallel A_2$ parallel agent in Example 4.

Example 4 (Watchdog)

In this example let

$$A_1 :: \text{do } (\text{tell}(c_1) \rightarrow^{+\infty} \text{ask}(c_3) \rightarrow^{+\infty} \text{success}) \text{ watching}^{+\infty}(c_2) \text{ else} \\ (\text{tell}(c_3) \rightarrow^{+\infty} \text{success})$$

and

$$A_2 :: \text{tell}(c_2) \rightarrow^{+\infty} \text{success}.$$

We evaluate the following watchdog construct with two agents A_1 and A_2 in parallel:

$$\langle (\text{do } (\text{tell}(c_1) \rightarrow^{+\infty} \text{ask}(c_3) \rightarrow^{+\infty} \text{success}) \text{ watching}^{+\infty}(c_2) \text{ else} \\ (\text{tell}(c_3) \rightarrow^{+\infty} \text{success}) \parallel \text{tell}(c_2) \rightarrow^{+\infty} \text{success}), \bar{0} \rangle.$$

According to Figure 3, agent A_1 is translated in the following way, where the agent B is a shorthand for the “else” branch of the watchdog, that is $\text{tell}(c_3) \rightarrow^{+\infty} \text{success}$:

$$\text{now}^{+\infty} c_2 \text{ then } B \text{ else } (\text{tell}(c_1) \rightarrow^{+\infty} \text{now}^{+\infty} c_2 \text{ then } B \text{ else} \\ (\text{ask}(c_3) \rightarrow^{+\infty} \text{now}^{+\infty} c_2 \text{ then } B \text{ else success})).$$

The execution timeline for this parallel agent is shown in Figure 7. In the first time-unit we have that $\sigma = \bar{0} \not\sqsubseteq c_2$, i.e., the store does not imply the guard of the $\text{now}^{+\infty}$, and therefore the interruption of the watchdog in A_1 is not triggered yet. Thus, in the first time-unit, both $\text{tell}(c_1) \rightarrow^{+\infty}$ of agent A_1 and $\text{tell}(c_2) \rightarrow^{+\infty}$ of agent A_2 are executed. At time $t = 1$, the interruption of the watchdog is immediately activated (i.e. $\text{now}^{+\infty} c_2$), since the store is now equal to $c_1 \otimes c_2 = c_3$ and $c_3 \vdash c_2$ (rule **R8** in Figure 2). Therefore, $\text{tell}(c_3) \rightarrow^{+\infty}$ of agent B in A_1 is executed, while A_2 already corresponds to the **success** agent).

5.1 An Auction Example

In Figure 8 we model the negotiation and the management of a generic service offered with a sort of auction: auctions, as other forms of negotiation, naturally need both timed and quantitative means to describe the interactions among agents. We reckon that an auction provides one of the most suitable example where to show the expressivity of the *tscpp* language, since both time and preference (for a service or

an object) are considered. In the following of the description we consider a *buyout* auction (Gallien and Gupta 2007), where the auctioneer improves the service and the related consumed resources (or, alternatively, its money price), bid after bid. When one (ore more) of the bidders agrees with the offer, it bids for it and the auction is immediately declared as over.

The auctioneer (i.e. *AUCTIONEER* in Figure 8) begins by offering a service described with the soft constraint c_{A_1} . We suppose that the cost associated to the soft constraint is expressed in terms of computational capabilities needed to support the execution of the service: e.g., $c_i \sqsubseteq c_j$ means that the service described by c_i needs more computational resources than c_j . By choosing the proper semiring, this load can be expressed as a percentage of the CPU use, or in terms of money, for example; we left this preference generic in the example, since we focus on the interaction among the agents.

We suppose that a constraint can be defined over three domains of QoS features: availability, reliability and execution time. For instance, c_{A_1} is defined as $availability > 95\% \wedge reliability > 99\% \wedge execution\ time < 3sec$. Clearly, providing a higher availability or reliability, and a lower execution time implies raising the computational resources to support this improvement, thus worsening the preference of the store.

```

AUCTIONEER ::
INIT_A  $\longrightarrow$ 
tell( $c_{A_1}$ )  $\xrightarrow{t_{sell}}$  ( $\sum_{i=1}^n ask(bidder_i = i) \rightarrow^{a_A} tell(winner = i) \rightarrow CHECK$ ) timeout( $w_A$ )
(tell( $c_{A_2}$ )  $\xrightarrow{t_{sell}}$  ( $\sum_{i=1}^n ask(bidder_i = i) \rightarrow^{a_A} tell(winner = i) \rightarrow CHECK$ ) timeout( $w_A$ )
(tell( $c_{A_3}$ )  $\xrightarrow{t_{sell}}$  ( $\sum_{i=1}^n ask(bidder_i = i) \rightarrow^{a_A} tell(winner = i) \rightarrow CHECK$ ) timeout( $w_A$ )
success))

CHECK ::
do ( ask(service = end)  $\longrightarrow$  success ) timeout( $w_C$ ) tell(service = interrupt) )
  watching $_{\phi_{check}}$ ( $c_{check}$ ) else (tell(service = interrupt)  $\longrightarrow$  STOP $_C$ )

BIDDER $_i$  ::
INIT_B $_i$   $\longrightarrow$ 
do ( TASK $_i$  ) watching $_{\phi_{bidder}}$ ( $c_{B_i}$ ) else ask( $\bar{1}$ )  $\xrightarrow{t_{buy_i}}$  tell(bidder $_i = i$ )  $\longrightarrow$ 
  ( ask(winner = i)  $\longrightarrow$  USER $_i$ ) + (ask(winner  $\neq$  i)  $\longrightarrow$  success) )

USER $_i$  ::
do ( USE_SERVICE $_i$   $\longrightarrow$  tell(service = end)  $\longrightarrow$  success )
  watching $_{\phi_{user}}$ (service = interrupt) else (STOP $_i$ )

AUCTION&MONITOR :: AUCTIONEER || BIDDER $_1$  || BIDDER $_2$  || ... || BIDDER $_n$ 

```

Fig. 8. An “auction and management” example for a generic service

After the offer, the auctioneer gives time to the bidders (each of them described with a possibly different agent *BIDDER $_i$* in Figure 8) to make their offer, since the choice of the winner is delayed by t_{sell} time-units (as in many real-world auction schemes). A level a_A is used to effectively check that the global consistency of the store is enough good, i.e., the computational power would not be already consumed under the given threshold. After the winner is nondeterministically chosen among all the bidders asking for the service, the auctioneer becomes a supervisor of the

used resource by executing the agent *CHECK*. Otherwise, if no offer is received within w_A time-units, a timeout interrupts the wait and the auctioneer improves the offered service by adding a new constraint: for example, in **tell**(c_{A_2}), c_{A_2} could be equivalent to *execution time* $< 1sec$, thus reducing the latency of the service (from 3 to 1 second) and consequently raising, at the same time, its computational cost (i.e., $\sigma = c_{A_1} \otimes c_{A_2} \sqsubseteq c_{A_1}$ means that we worsen the consistency level of the store). The same offer/wait process is repeated three times in Figure 8.

Each of the bidders in Figure 8 executes its own task (i.e., $TASK_i$, left generic since not in the scope of the example), but as soon as the offered resource meets its demand (i.e. c_{B_i} is satisfied by the store: $\sigma \sqsubseteq c_{B_i}$), the bidder is interrupted and then asks to use the service. The time needed to react and make an offer is modeled with t_{buy_i} : fast bidders will have more chances to win the auction, if their request arrives before the choice of the auctioneer. If one of the bidders wins, then it becomes a user of the resource, by executing *USER_i*.

The agent *USER_i* uses the service (through the agent *USE_SERVICE_i*, left generic in Figure 8), but it stops (using agent *STOP_i*, left generic in Figure 8) as soon as the service is interrupted, i.e., as the store satisfies *service* = *interrupt*. On the other side, agent *CHECK* waits for the use termination, but it interrupts the user if the computation takes too long (more than w_C time-units), or if the user absorbs the computational capabilities beyond a given threshold, i.e. as soon as the c_{check} becomes implied by the store (i.e. $\sigma \sqsubseteq c_{check}$): in fact, *USE_SERVICE_i* could be allowed to ask for more power by “telling” some more constraints to the store. To interrupt the service use, agent *CHECK* performs a **tell**(*service* = *interrupt*). All the agents *INIT*, left generic in Figure 8, can be used to initialize the computation. In order to avoid a heavy notation in Figure 8, we do not show the preference associated to constraints and the consistency check label on the transition arrows, when they are not significative for the example description. Also the ϕ_{Check} , ϕ_{Bidder} and ϕ_{User} thresholds of the watchdog constructs are not detailed.

Finally, in the following we model a more refined behaviour of the auctioneer, which accepts the bidding with the highest value, where *CHECK*, *BIDDER_i* and *USER_i* are defined as in Figure 8.

```

AUCTIONEER' ::
INIT_A →
tell( $c_{A_1}$ )  $\xrightarrow{t_{sell}}$  ( $\sum_{i=1}^n \text{ask}(bidder_i = i) \rightarrow^{a_A} \text{CHOOSE}$ ) timeout( $w_A$ )
( $\text{tell}(c_{A_2}) \xrightarrow{t_{sell}}$  ( $\sum_{i=1}^n \text{ask}(bidder_i = i) \rightarrow^{a_A} \text{CHOOSE}$ ) timeout( $w_A$ ))
( $\text{tell}(c_{A_3}) \xrightarrow{t_{sell}}$  ( $\sum_{i=1}^n \text{ask}(bidder_i = i) \rightarrow^{a_A} \text{CHOOSE}$ ) timeout( $w_A$ ) success))

CHOOSE ::
now ( $bidder_n = n$ ) then tell( $winner = n$ ) → CHECK
else ( now ( $bidder_{n-1} = n-1$ ) then tell( $winner = n-1$ ) → CHECK
      else ( ... ( now ( $bidder_2 = 2$ ) then tell( $winner = 2$ ) → CHECK
                  else tell( $winner = 1$ ) → CHECK) ... ) )

newAUCTION&MONITOR :: AUCTIONEER' || BIDDER1 || BIDDER2 || ... || BIDDERn

```

Fig. 9. A new “auction and management” example for a generic service

Many other real-life automated tasks can be modeled with the *tscpp* language.

For example, a quality-driven composition of web services: the agents that represent different web services can add to the store their functionalities (represented by soft constraints) with **tell** actions; the final store models their composition. The consistency level of the store represents (for example) the total monetary cost of the obtained service, or a value representing the consistency of the integrated functionalities. The reason is that, when we compose the services offered by different providers, we cannot be sure of how much they are compatible. A client wishing to use the composed service can perform an **ask** with a threshold such that it prevents the client from paying a high price, or having an unreliable service. Softness is also useful to model incomplete service specifications that may evolve incrementally and, in general, for non-functional aspects.

6 The Denotational Model

In this section we define a denotational characterization of the operational semantics obtained by following the construction in (de Boer et al. 2000), and by using *timed reactive sequences* to represent *tsccp* computations. These sequences are similar to those used in the semantics of dataflow languages (Jonsson 1985), imperative languages (Brookes 1993) and (timed) *ccp* (de Boer and Palamidessi 1991; de Boer et al. 2000).

The denotational model associates with a process a set of timed reactive sequences of the form $\langle \sigma_1, \gamma_1 \rangle \cdots \langle \sigma_n, \gamma_n \rangle \langle \sigma, \sigma \rangle$ where a pair of constraints $\langle \sigma_i, \gamma_i \rangle$ represents a reaction of the given process at time i : intuitively, the process transforms the global store from σ_i to γ_i or, in other words, σ_i is the assumption on the external environment while γ_i is the contribution of the process itself (which always entails the assumption). The last pair denotes a “stuttering step” in which the agent **Success** has been reached. Since the basic actions of *tsccp* are monotonic and we can also model a new input of the external environment by a corresponding tell operation, it is natural to assume that reactive sequences are monotonic. Thus, in the following we assume that each timed reactive sequence $\langle \sigma_1, \gamma_1 \rangle \cdots \langle \sigma_{n-1}, \gamma_{n-1} \rangle \langle \sigma_n, \sigma_n \rangle$ satisfies the conditions $\gamma_i \vdash \sigma_i$ and $\sigma_j \vdash \gamma_{j-1}$, for any $i \in [1, n-1]$ and $j \in [2, n]$.

The set of all reactive sequences is denoted by \mathcal{S} , its typical elements by $s, s_1 \dots$, while sets of reactive sequences are denoted by $S, S_1 \dots$, and ε indicates the empty reactive sequence. Furthermore, the symbol \cdot denotes the operator that concatenates sequences. In the following, *Process* denotes the set of *tsccp* processes.

Operationally, the reactive sequences of an agent are generated as follows.

Definition 3 (Processes Semantics)

We define the semantics $\mathcal{R} \in \text{Process} \rightarrow \mathcal{P}(\mathcal{S})$ by

$$\begin{aligned} \mathcal{R}(F.A) = & \{ \langle \sigma, \sigma' \rangle \cdot w \in \mathcal{S} \mid \langle A, \sigma \rangle \longrightarrow \langle B, \sigma' \rangle \text{ and } w \in \mathcal{R}(F.B) \} \\ & \cup \\ & \{ \langle \sigma, \sigma \rangle \cdot w \in \mathcal{S} \mid \langle A, \sigma \rangle \not\rightarrow \text{ and } \\ & \quad \text{either } A \neq \mathbf{Success} \text{ and } w \in \mathcal{R}(F.A) \\ & \quad \text{or } A = \mathbf{Success} \text{ and } w \in \mathcal{R}(F.A) \cup \{\varepsilon\} \}. \end{aligned}$$

Formally \mathcal{R} is defined as the least fixed-point of the operator $\Phi \in (Process \rightarrow \mathcal{P}(\mathcal{S})) \rightarrow Process \rightarrow \mathcal{P}(\mathcal{S})$ defined by

$$\begin{aligned} \Phi(I)(F.A) = & \{ \langle \sigma, \delta \rangle \cdot w \in \mathcal{S} \mid \langle A, \sigma \rangle \longrightarrow \langle B, \delta \rangle \text{ and } w \in I(F.B) \} \\ & \cup \\ & \{ \langle \sigma, \sigma \rangle \cdot w \in \mathcal{S} \mid \langle A, \sigma \rangle \not\rightarrow \text{ and } \\ & \quad \text{either } A \neq \mathbf{Success} \text{ and } w \in I(F.A) \\ & \quad \text{or } A = \mathbf{Success} \text{ and } w \in I(F.A) \cup \{\varepsilon\} \}. \end{aligned}$$

The ordering on $Process \rightarrow \mathcal{P}(\mathcal{S})$ is that of (point-wise extended) set-inclusion, and since it is straightforward to check that Φ is continuous, standard results ensure that the least fixpoint exists (and it is equal to $\sqcup_{n \geq 0} \Phi^n(\perp)$).

Note that $\mathcal{R}(F.A)$ is the union of the set of all successful reactive sequences that start with a reaction of A , and the set of all successful reactive sequences that start with a stuttering step of A . In fact, when an agent is blocked, i.e., it cannot react to the input of the environment, a stuttering step is generated. After such a stuttering step, the computation can either continue with the further evaluation of A (possibly generating more stuttering steps), or it can terminate if A is the **Success** agent. Note also that, since the **Success** agent used in the transition system cannot make any move, an arbitrary (finite) sequence of stuttering steps is always appended to each reactive sequence.

6.1 Correctness

The observables $\mathcal{O}_{io}^{mp}(P)$ describing the input/output pairs of successful computations can be obtained from $\mathcal{R}(P)$ by considering suitable sequences, namely those sequences which do not perform assumptions on the store. In fact, note that some reactive sequences do not correspond to real computations: Clearly, when considering a real computation no further contribution from the environment is possible. This means that, at each step, the assumption on the current store must be equal to the store produced by the previous step. In other words, for any two consecutive steps $\langle \sigma_i, \sigma'_i \rangle \langle \sigma_{i+1}, \sigma'_{i+1} \rangle$ we must have $\sigma'_i = \sigma_{i+1}$. Thus, we are led to the following.

Definition 4 (Connected Sequences)

Let $s = \langle \sigma_1, \sigma'_1 \rangle \langle \sigma_2, \sigma'_2 \rangle \cdots \langle \sigma_n, \sigma'_n \rangle$ be a reactive sequence. We say that s is connected if $\sigma_1 = \mathbf{1}$ and $\sigma_i = \sigma'_{i-1}$ for each i , $2 \leq i \leq n$.

According to the previous definition, a sequence is connected if all the information assumed on the store is produced by the process itself. To be defined as connected, a sequence must also have $\mathbf{1}$ as the initial constraint. A connected sequence $s = \langle \mathbf{1}, \sigma_1 \rangle \langle \sigma_1, \sigma_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle$ represents a *tsccp* computation of a process $F.A$, where $\mathbf{1}$ is the input constraint and $\sigma_n \Downarrow_{Fv(A)}$ is the result. From the above discussion we can derive the following property:

Proposition 1 (Correctness)

For any process $P = F.A$ we have

$$\mathcal{O}_{io}^{mp}(P) = \{ \sigma_n \Downarrow_{Fv(A)} \mid \text{there exists a connected sequence } s \in \mathcal{R}(P) \text{ such that } s = \langle \mathbf{1}, \sigma_1 \rangle \langle \sigma_1, \sigma_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle \}.$$

Proof

From the close correspondence between the rules of the transition system and the definition of the denotational semantics, we have that $s \in \mathcal{R}(P)$ if and only if $s = \langle \sigma_1, \sigma'_1 \rangle \langle \sigma_2, \sigma'_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle$, $A_1 = A$, $A_n = \mathbf{Success}$ and for $i \in [1, n-1]$,

- either $\langle A_i, \sigma_i \rangle \longrightarrow \langle A_{i+1}, \sigma'_i \rangle$
- or $\langle A_i, \sigma_i \rangle \not\longrightarrow$, $A_{i+1} = A_i$ and $\sigma'_i = \sigma_i$.

Then there exists a connected sequence $s \in \mathcal{R}(P)$ if and only if $s = \langle \sigma_1, \sigma_2 \rangle \langle \sigma_2, \sigma_3 \rangle \cdots \langle \sigma_n, \sigma_n \rangle$, $A_1 = A$, $\sigma_1 = \mathbf{1}$, $A_n = \mathbf{Success}$ and for $i \in [1, n-1]$, $\langle A_i, \sigma_i \rangle \longrightarrow \langle A_{i+1}, \sigma_{i+1} \rangle$. Therefore, the proof follows by definition of $\mathcal{O}_{io}^{mp}(P)$. \square

6.2 Compositionality of the Denotational Semantics for *tsccp* Processes

In order to prove the compositionality of the denotational semantics, we now introduce a semantics $\llbracket F.A \rrbracket(e)$, which is compositional by definition and where, for technical reasons, we explicitly represent the environment e that associates a denotation to each procedure identifier. More precisely, assuming that $Pvar$ denotes the set of procedure identifiers, $Env = Pvar \rightarrow \mathcal{P}(\mathcal{S})$, with typical element e , is the set of *environments*. Given $e \in Env$, $p \in Pvar$ and $f \in \mathcal{P}(\mathcal{S})$, we denote by $e' = e\{f/p\}$ the new environment such that $e'(p) = f$ and $e'(p') = e(p')$ for each procedure identifier $p' \neq p$.

Given a process $F.A$, the denotational semantics $\llbracket F.A \rrbracket : Env \rightarrow \mathcal{P}(\mathcal{S})$ is defined by the equations in Figure 10, where μ denotes the least fixpoint with respect to the subset inclusion of elements of $\mathcal{P}(\mathcal{S})$. The semantic operators appearing in Figure 10 are formally defined as follows; intuitively they reflect the operational behavior of their syntactic counterparts in terms of reactive sequences.⁵ We first need the following definition.

Definition 5

Let σ, ϕ and c be constraints in \mathcal{C} and let $a \in \mathcal{A}$. We say that

- $\sigma \succ^a c$, if $(\sigma \vdash c \text{ and } \sigma \Downarrow_\emptyset \not\prec a)$ while $\sigma \succ_\phi c$, if $(\sigma \vdash c \text{ and } \sigma \not\vdash \phi)$.

Definition 6 (Semantic operators)

Let S, S_i be sets of reactive sequences, c, c_i be constraints and let \succ_i be either of the form \succ^{a_i} or \succ_{ϕ_i} . Then we define the operators $t\ell l$, \sum , \parallel , $n\tilde{o}w$ and $\exists x$ as follows:

The (valued) tell operator

$$t\tilde{\ell}l^a(c, S) = \{s \in \mathcal{S} \mid s = \langle \sigma, \sigma \otimes c \rangle \cdot s', \sigma \otimes c \Downarrow_\emptyset \not\prec a \text{ and } s' \in S\}.$$

$$t\tilde{\ell}l_\phi(c, S) = \{s \in \mathcal{S} \mid s = \langle \sigma, \sigma \otimes c \rangle \cdot s', \sigma \otimes c \not\vdash \phi \text{ and } s' \in S\}.$$

⁵ In Figure 10 the syntactic operator \rightarrow_i is either of the form \rightarrow^{a_i} or \rightarrow_{ϕ_i} .

The guarded choice

$$\sum_{i=1}^n c_i \succ_i S_i = \{ s \cdot s' \in \mathcal{S} \mid \begin{array}{l} s = \langle \sigma_1, \sigma_1 \rangle \cdots \langle \sigma_m, \sigma_m \rangle, \sigma_j \not\prec_i c_i \\ \text{for each } j \in [1, m-1], i \in [1, n], \\ \sigma_m \succ_h c_h \text{ and } s' \in S_h \text{ for an } h \in [1, n] \end{array} \}.$$

The parallel composition Let $\tilde{\parallel} \in \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S}$ be the (commutative and associative) partial operator defined as follows:

$$\langle \sigma_1, \sigma_1 \otimes \gamma_1 \rangle \cdots \langle \sigma_n, \sigma_n \otimes \gamma_n \rangle \langle \sigma, \sigma \rangle \tilde{\parallel} \langle \sigma_1, \sigma_1 \otimes \delta_1 \rangle \cdots \langle \sigma_n, \sigma_n \otimes \delta_n \rangle \langle \sigma, \sigma \rangle = \langle \sigma_1, \sigma_1 \otimes \gamma_1 \otimes \delta_1 \rangle \cdots \langle \sigma_n, \sigma_n \otimes \gamma_n \otimes \delta_n \rangle \langle \sigma, \sigma \rangle.$$

We define $S_1 \tilde{\parallel} S_2$ as the point-wise extension of the above operator to sets.

The (valued) now operator

$$now^a(c, S_1, S_2) = \{ s \in \mathcal{S} \mid \begin{array}{l} s = \langle \sigma, \sigma' \rangle \cdot s', \sigma \Downarrow_{\emptyset} \not\prec a \text{ and} \\ \text{either } \sigma \vdash c \text{ and } s \in S_1 \\ \text{or } \sigma \not\vdash c \text{ and } s \in S_2 \end{array} \}.$$

$$now_\phi(c, S_1, S_2) = \{ s \in \mathcal{S} \mid \begin{array}{l} s = \langle \sigma, \sigma' \rangle \cdot s', \sigma \not\vdash \phi \text{ and} \\ \text{either } \sigma \vdash c \text{ and } s \in S_1 \\ \text{or } \sigma \not\vdash c \text{ and } s \in S_2 \end{array} \}.$$

The hiding operator The semantic hiding operator can be defined as follows:

$$\tilde{\exists}xS = \{ s \in \mathcal{S} \mid \text{there exists } s' \in S \text{ such that } s = s'[x/y] \text{ with } y \text{ new} \}$$

where $s'[x/y]$ denotes the sequence obtained from s' by replacing the variable x for the variable y , which we assume to be new.⁶

Obviously, the semantic (valued) tell operator reflects the operational behavior of the syntactic (valued) tell. Concerning the semantic choice operator, a sequence in $\sum_{i=1}^n c_i \succ_i S_i$ consists of an initial period of waiting for a store which satisfies one of the guards. During this waiting period, only the environment is active by producing the constraints σ_j , while the process itself generates the stuttering steps $\langle \sigma_j, \sigma_j \rangle$. When the store is strong enough to satisfy a guard, that is to entail a c_h and to satisfy the condition on the cut level, the resulting sequence is obtained by adding $s' \in S_h$ to the initial waiting period. In the semantic parallel operator defined on sequences, we require that the two arguments of the operator agree at each point of time with respect to the contribution of the environment (the σ_i 's), and that they have the same length (in all other cases the parallel composition is assumed being undefined).

If $F.A$ is a closed process, that is if all the procedure names occurring in A are defined in F , then $\llbracket F.A \rrbracket(e)$ does not depend on e , and it will be indicated as $\llbracket F.A \rrbracket$. Environments in general allow us to define the semantics also of processes that are not closed. The following result shows the correspondence between the two semantics we have introduced and, therefore, it proves the compositionality of $\mathcal{R}(F.A)$. From the above discussion we can derive the following property:

⁶ To be more precise, we assume that each time that we consider a new application of the operator $\tilde{\exists}$ we use a new, different y . As in the case of the operational semantics, this can be ensured by a suitable renaming mechanism.

Proposition 2 (Compositionality)

If $F.A$ is closed then $\mathcal{R}(F.A) = \llbracket F.A \rrbracket$ holds.

Proof

We prove by induction on the complexity of the agent A that

$$\begin{aligned} \llbracket F.A \rrbracket = \{s \mid & s = \langle \sigma_1, \sigma'_1 \rangle \langle \sigma_2, \sigma'_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle, \\ & A_1 = A, A_n = \mathbf{Success} \text{ and for } i \in [1, n-1], \\ & \text{either } \langle A_i, \sigma_i \rangle \longrightarrow \langle A_{i+1}, \sigma'_i \rangle \\ & \text{or } \langle A_i, \sigma_i \rangle \not\longrightarrow, A_{i+1} = A_i, \sigma'_i = \sigma_i \}. \end{aligned}$$

Then the proof follows by definition of $\mathcal{R}(P)$.

When the P is not of the form $F.B \parallel C$ the thesis follows immediately from the close correspondence between the rules of the transition system and the definition of the denotational semantics.

Assume now that P is of the form $F.B \parallel C$. By definition of the denotational semantics, $s \in \llbracket F.A \rrbracket$ if and only if $s = \langle \sigma_1, \sigma'_1 \rangle \langle \sigma_2, \sigma'_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle$ and there exist $s' \in \llbracket F.B \rrbracket$ and $s'' \in \llbracket F.C \rrbracket$,

$$\begin{aligned} s' &= \langle \sigma_1, \sigma_1 \otimes \gamma_1 \rangle \langle \sigma_2, \sigma_2 \otimes \gamma_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle \\ s'' &= \langle \sigma_1, \sigma_1 \otimes \delta_1 \rangle \langle \sigma_2, \sigma_2 \otimes \delta_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle \end{aligned}$$

such that for each $i \in [1, n-1]$, $\sigma'_i = \sigma_i \otimes \gamma_i \otimes \delta_i$. By inductive hypothesis $s' \in \llbracket F.B \rrbracket$ and $s'' \in \llbracket F.C \rrbracket$ if and only if for $i \in [1, n-1]$,

$$\begin{aligned} & \text{either } \langle B_i, \sigma_i \rangle \longrightarrow \langle B_{i+1}, \sigma_i \otimes \gamma_i \rangle, \\ & \text{or } \langle B_i, \sigma_i \rangle \not\longrightarrow, B_{i+1} = B_i, \sigma_i \otimes \gamma_i = \sigma_i \quad \text{and} \\ & \text{either } \langle C_i, \sigma_i \rangle \longrightarrow \langle C_{i+1}, \sigma_i \otimes \delta_i \rangle, \\ & \text{or } \langle C_i, \sigma_i \rangle \not\longrightarrow, C_{i+1} = C_i, \sigma_i \otimes \delta_i = \sigma_i. \end{aligned} \tag{1}$$

$B_1 = B$, $B_n = \mathbf{Success}$, $C_1 = C$ and $C_n = \mathbf{Success}$. Therefore, by Rule **R8** and previous observations, we have that (1) holds if and only if $B_1 \parallel C_1 = B \parallel C$, $B_n \parallel C_n = \mathbf{Success}$ and for $i \in [1, n-1]$,

$$\begin{aligned} & \text{either } \langle B_i \parallel C_i, \sigma_i \rangle \longrightarrow \langle B_{i+1} \parallel C_{i+1}, \sigma'_i \rangle \\ & \text{or } \langle B_i \parallel C_i, \sigma_i \rangle \not\longrightarrow, A_{i+1} \parallel B_{i+1} = A_i \parallel B_i, \sigma'_i = \sigma_i \end{aligned}$$

and then the thesis. \square

7 An Interleaving Approach for non-Time-elapsing Actions

In this section, we show a different version of the *tsccp* language: while in *tsccp* the parallel operator is modeled in terms of *maximal parallelism*, the same operator can be treated also in terms of interleaving. According to *maximal parallelism*, at each moment every enabled agent of the system is activated, while in the second paradigm an agent could not be assigned to a “free” processor. Clearly, since we have dynamic process creation, a maximal parallelism approach has the disadvantage that, in general, it implies the existence of an unbound number of processes. On the other hand a naive interleaving semantic could be problematic from the

- E1** $\llbracket F.\text{success} \rrbracket(e) = \{\langle \sigma_1, \sigma_1 \rangle \langle \sigma_2, \sigma_2 \rangle \cdots \langle \sigma_n, \sigma_n \rangle \in \mathcal{S} \mid n \geq 1\}$
- E2** $\llbracket F.\text{tell}(c) \rightarrow^a A \rrbracket(e) = \tilde{\text{tell}}^a(c, \llbracket F.A \rrbracket(e))$
- E3** $\llbracket F.\text{tell}(c) \rightarrow_\phi A \rrbracket(e) = \tilde{\text{tell}}_\phi(c, \llbracket F.A \rrbracket(e))$
- E4** $\llbracket F.\sum_{i=1}^n \text{ask}(c_i) \rightarrow_i A_i \rrbracket(e) = \tilde{\sum}_{i=1}^n c_i \succ_i \llbracket F.A_i \rrbracket(e)$
- E5** $\llbracket F.\text{now}^a c \text{ then } A \text{ else } B \rrbracket(e) = n\tilde{\text{ow}}^a(c, \llbracket F.A \rrbracket(e), \llbracket F.B \rrbracket(e))$
- E6** $\llbracket F.\text{now}_\phi c \text{ then } A \text{ else } B \rrbracket(e) = n\tilde{\text{ow}}_\phi(c, \llbracket F.A \rrbracket(e), \llbracket F.B \rrbracket(e))$
- E7** $\llbracket F.A \parallel B \rrbracket(e) = \llbracket F.A \rrbracket(e) \parallel \llbracket F.B \rrbracket(e)$
- E8** $\llbracket F.\exists x A \rrbracket(e) = \tilde{\exists}x \llbracket F.A \rrbracket(e)$
- E9** $\llbracket F.p(x) \rrbracket(e) = \mu\Psi \text{ where } \Psi(f) = \llbracket F \setminus \{p\}.\text{ask}(\bar{1}) \rightarrow A \rrbracket(e\{f/p\}), \quad p(x) :: A \in F$

Fig. 10. The semantics $\llbracket F.A \rrbracket(e)$.

time viewpoint, as in principle the time does not pass for enabled agent which are not scheduled. For the semantics in this section we follow a solution analogous to that one adopted in (de Boer et al. 2004): we assume that the parallel operator is interpreted in terms of interleaving, as usual, however we must assume maximal parallelism for actions depending on time. In other words, time passes for all the parallel processes involved in a computation. To summarize, in this section we adopt maximal parallelism for time elapsing (i.e. for timeout constructs) and an interleaving model for basic computation steps (i.e. (valued) **ask** and (valued) **tell** actions).

To distinguish this new approach, we named the resulting language as *tsccp-i*, i.e., *tsccp* with interleaving. Time-outs are modeled in *tsccp-i* by the construct **askp**_{*t*}(*c*)?_{*φ*}*A*:*B* which replaces the **now**_{*φ*} *c* **then** *A* **else** *B* construct of *tsccp* and directly has time *t* as one of its parameters, differently from the **now**_{*φ*} agent. The **askp**_{*t*} agent can be interpreted as follows: one is allowed to wait *t* time-units for the entailment of the constraint *c* by the store and the subsequent evaluation of the process *A*; if this time limit is exceeded, then the process *B* is evaluated. Analogously for the construct **askp**_{*t*}(*c*)?^{*a*}*A*:*B*.

Definition 7 (*tsccp-i*)

Given a soft constraint system $\langle S, D, V \rangle$, the corresponding structure \mathcal{C} , any semir-

ing value a , *soft constraints* $\phi, c \in \mathcal{C}$ and any tuple of variables x , the syntax of the *tsccp-i* language is given by the following grammar:

$$\begin{aligned}
P &::= F.A \\
F &::= p(x) :: A \mid F \cdot F \\
A &::= \text{success} \mid \text{tell}(c) \rightarrow_{\phi} A \mid \text{tell}(c) \rightarrow^a A \mid E \mid A \parallel A \mid \exists x A \mid p(x) \mid \\
&\quad \Sigma_{i=1}^n E_i \mid \text{askp}_t(c)?_{\phi} A:A \mid \text{askp}_t(c)?^a A:A \\
E &::= \text{ask}(c) \rightarrow_{\phi} A \mid \text{ask}(c) \rightarrow^a A
\end{aligned}$$

where, as in Definition 1, P is the class of processes, F is the class of sequences of procedure declarations (or clauses), A is the class of agents. As before, in a *tsccp-i* process $P = F.A$, A is the initial agent, to be executed in the context of the set of declarations F .

Analogously to *tsccp* processes, in order to simplify the notation, in the following we will usually write a *tsccp-i* process $P = F.A$ simply as the corresponding agent A .

The operational model of *tsccp-i* processes can be formally described by a labeled transition system $T = (\text{Conf}, \text{Label}, \longrightarrow)$, where we assume that each transition step exactly takes one time-unit. Configurations (in) Conf are pairs consisting of a process and a constraint in \mathcal{C} representing the common *store*. $\mathcal{L} = \{\tau, \omega\}$ is the set of labels. We use labels to distinguish “real” computational steps performed by processes which have the control (label ω) from the transitions which model only the passing of time (label τ). So ω -actions are those performed by processes that modify the store (**tell**), perform a check on the store (**ask**, **askp_t**), correspond to exceeding a time-out (**askp₀**), or perform a choice ($\Sigma_{i=1}^n E_i$). On the other hand, τ -actions are those performed by time-out processes (**askp_t**) in case they have not the control. In Figure 11 we show the semantics of all the *tsccp-i* actions, but in the following we describe only the actions whose semantics is different from that one presented in Figure 2 (i.e., for *tsccp*), that is we describe in detail the parallelism and the **askp_t** agent. The semantics of the other actions of *tsccp-i* is the same as for *tsccp*, except for the fact that their transition is labeled with ω .

Parallelism Rules **Q5** and **Q6** in Figure 11 model the parallel composition operator in terms of *interleaving*, since only one basic ω -action is allowed for each transition (i.e. for each unit of time). This means that the access to the shared store is granted to one process a time. However, time passes for all the processes appearing in the \parallel context at the external level, as shown by rule **Q5**, since τ -actions are allowed together with a ω -action. On the other hand, a parallel component is allowed to proceed in isolation if (and only if) the other parallel component cannot perform a τ -action (rule **Q6**). To summarize, we adopt maximal parallelism for time elapsing (i.e. τ -actions) and an interleaving model for basic computation steps (i.e. ω -actions).

We have adopted this approach because it seems more adequate to the nature of time-out operators not to interrupt the elapsing of time, once the evaluation of a time-out has started. Clearly one could start the elapsing of time when the time out process is scheduled, rather than when it appears in the top-level current

Q1	$\frac{(\sigma \otimes c) \Downarrow_{\emptyset} \not\prec a}{\langle \text{tell}(c) \rightarrow^a A, \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \otimes c \rangle}$	V-tell
Q2	$\frac{\sigma \otimes c \not\sqsubseteq \phi}{\langle \text{tell}(c) \rightarrow_{\phi} A, \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \otimes c \rangle}$	Tell
Q3	$\frac{\sigma \vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a}{\langle \text{ask}(c) \rightarrow^a A, \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \rangle}$	V-ask
Q4	$\frac{\sigma \vdash c \quad \sigma \not\sqsubseteq \phi}{\langle \text{ask}(c) \rightarrow_{\phi} A, \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \rangle}$	Ask
Q5	$\frac{\langle A, \sigma \rangle \xrightarrow{\xi} \langle A', \sigma' \rangle \quad \langle B, \sigma \rangle \xrightarrow{\tau} \langle B', \sigma \rangle \quad \xi \in \{\tau, \omega\}}{\langle A \parallel B, \sigma \rangle \xrightarrow{\xi} \langle A' \parallel B', \sigma' \rangle}$ $\langle B \parallel A, \sigma \rangle \xrightarrow{\xi} \langle B' \parallel A', \sigma' \rangle$	Parall1
Q6	$\frac{\langle A, \sigma \rangle \xrightarrow{\xi} \langle A', \sigma' \rangle \quad \langle B, \sigma \rangle \not\xrightarrow{\tau} \quad \xi \in \{\tau, \omega\}}{\langle A \parallel B, \sigma \rangle \xrightarrow{\xi} \langle A' \parallel B, \sigma' \rangle}$ $\langle B \parallel A, \sigma \rangle \xrightarrow{\xi} \langle B \parallel A', \sigma' \rangle$	Parall2
Q7	$\frac{\langle E_j, \sigma \rangle \xrightarrow{\omega} \langle A_j, \sigma' \rangle \quad j \in [1, n]}{\langle \Sigma_{i=1}^n E_i, \sigma \rangle \xrightarrow{\omega} \langle A_j, \sigma' \rangle}$	Nondet
Q8	$\langle p(x), \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \rangle \quad p(x) :: A \in F$	P-call
Q9	$\frac{\langle A[x/y], \sigma \rangle \xrightarrow{\xi} \langle B, \sigma' \rangle \quad \xi \in \{\tau, \omega\}}{\langle \exists x A, \sigma \rangle \xrightarrow{\xi} \langle B, \sigma' \rangle}$	Hide
Q10	$\frac{\sigma \vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a \quad t > 0}{\langle \text{askp}_t(c)?^a A:B, \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \rangle}$	V-askp1
Q11	$\frac{\sigma \Downarrow_{\emptyset} < a \quad t > 0}{\langle \text{askp}_t(c)?^a A:B, \sigma \rangle \xrightarrow{\omega} \langle B, \sigma \rangle}$	V-askp2
Q12	$\frac{\sigma \not\vdash c \quad \sigma \Downarrow_{\emptyset} \not\prec a \quad t > 0}{\langle \text{askp}_t(c)?^a A:B, \sigma \rangle \xrightarrow{\omega} \langle \text{askp}_{t-1}(c)?^a A:B, \sigma \rangle}$	V-askp3
Q13	$\langle \text{askp}_t(c)?^a A:B, \sigma \rangle \xrightarrow{\tau} \langle \text{askp}_{t-1}(c)?^a A:B, \sigma \rangle \quad t > 0$	V-askp4
Q14	$\langle \text{askp}_0(c)?^a A:B, \sigma \rangle \xrightarrow{\omega} \langle B, \sigma \rangle$	V-askp5
Q15	$\frac{\sigma \vdash c \quad \sigma \not\sqsubseteq \phi \quad t > 0}{\langle \text{askp}_t(c)?_{\phi} A:B, \sigma \rangle \xrightarrow{\omega} \langle A, \sigma \rangle}$	Askp1
Q16	$\frac{\sigma \sqsubseteq \phi \quad t > 0}{\langle \text{askp}_t(c)?_{\phi} A:B, \sigma \rangle \xrightarrow{\omega} \langle B, \sigma \rangle}$	Askp2
Q17	$\frac{\sigma \not\vdash c \quad \sigma \not\sqsubseteq \phi \quad t > 0}{\langle \text{askp}_t(c)?_{\phi} A:B, \sigma \rangle \xrightarrow{\omega} \langle \text{askp}_{t-1}(c)?_{\phi} A:B, \sigma \rangle}$	Askp3
Q18	$\langle \text{askp}_t(c)?_{\phi} A:B, \sigma \rangle \xrightarrow{\tau} \langle \text{askp}_{t-1}(c)?_{\phi} A:B, \sigma \rangle \quad t > 0$	Askp4
Q19	$\langle \text{askp}_0(c)?_{\phi} A:B, \sigma \rangle \xrightarrow{\omega} \langle B, \sigma \rangle$	Askp5

Fig. 11. The transition system for *tscpp-i*.

parallel context. This modification could easily be obtained by adding a syntactic construct to differentiate active timeouts from inactive ones, and by accordingly changing the transition system. One could also easily modify the semantics (both operational and denotational) to consider a more liberal assumption which allows multiple ask actions in parallel.

Valued-Askp_t The rules **Q10-Q14** in Figure 11 show that the time-out process $\text{askp}_t(c)?^a A:B$ behaves as A if c is entailed by the store and the store is “consistent enough” with respect to the threshold a in the next t time-units: if $t > 0$ and the condition on the store and the cut level are satisfied, then the agent A is evaluated (rule **Q10**). If $t > 0$ and the condition on the cut level is not satisfied, then the agent B is evaluated (rule **Q11**). Finally if $t > 0$, the condition on the cut level is satisfied, but the condition on the store is not satisfied, then the control is repeated at the next time instant and the value of the counter t is decreased (axiom **Q12**); note that in this case we use the label ω , since a check on the store has been performed. As shown by axiom **Q13**, the counter can be decreased also by performing a τ -action: intuitively, this rule is used to model the situation in which, even though the evaluation of the time-out started already, another (parallel) process has the control. In this case, analogously to the approach in (de Boer et al. 2004) and differently from the approach in (Busi et al. 2000), time continues to elapse (via τ -actions) also for the time-out process (see also the rules **Q5** and **Q6** of the parallel operator). Axiom **Q14** shows that, if the time-out is exceeded, i.e., the counter t has reached the value of 0, then the process $\text{askp}_t(c)?^a A:B$ behaves as B .

Askp_t The rules **Q15-Q19** in Figure 11 are similar to rules **Q10-Q14** described before, with the exception that here a finer (pointwise) threshold ϕ is compared to the store σ , analogously to what happens with the **tell** and **ask** agents.

In the following we provide the definition for the observables of the language, which are clearly based only on ω -actions.

Definition 8 (Observables for tsccp-i)

Let $P = F.A$ be a *tsccp-i* process. We define

$$\mathcal{O}_{io}^i(P) = \{\gamma \Downarrow_{Fv(A)} \mid \langle A, \bar{1} \rangle \xrightarrow{\omega}^* \langle \text{Success}, \gamma \rangle\},$$

where **Success** is any agent that contains only occurrences of the agent **success** and of the operator \parallel .

8 An Execution Timeline for a *tsccp-i* Parallel Agent

In this section we show a timeline for the execution of three *tsccp-i* agents in parallel. We consider the three soft constraints shown in Figure 4 and the *Weighted* semiring $\langle \mathbb{R}^+ \cup \{+\infty\}, \min, +, +\infty, 0 \rangle$ (Bistarelli 2004; Bistarelli et al. 1997). Our parallel agent is defined by:

$$\begin{aligned} A_1 &:: \text{askp}_5(c_3)?^{+\infty}(\text{tell}(c_1) \rightarrow^{+\infty} \text{success}):(\text{success}) \\ A_2 &:: \text{tell}(c_1) \rightarrow^{+\infty} \text{success} \\ A_3 &:: \text{tell}(c_2) \rightarrow^{+\infty} \text{success}. \end{aligned}$$

Their concurrent evaluation in the $\bar{0}$ empty store is shown in Figure 12. At $t = 0$ and $t = 1$ the agent A_1 can make a τ -transition (rule **Q13** in Figure 11), waiting for the elapsing of 1 time-unit. This can be done in parallel with a single other ω -action: therefore, the **tell**(c_1) of agent A_2 , and the **tell**(c_2) of agent A_3 cannot run in parallel at the same time, since they are both ω -actions. In the execution shown in Figure 12, A_2 is executed before A_3 (also the opposite is possible, depending on the scheduling), leading to the store $\sigma = c_1 \otimes c_2 = c_3$. At $t = 2$, the guard of **askp**₅ in agent A_1 is enabled since $\sigma \vdash c_3$ and, therefore, rule **Q10** in Figure 11 is executed. Finally, at $t = 3$ the **tell**(c_1) action of agent A_1 is executed as the last action, and at $t = 4$ we have $\langle \text{success} \parallel \text{success} \parallel \text{success}, c_1 \otimes c_2 \otimes c_1 \rangle$.

	$t=0$	$t=1$	$t=2$	$t=3$	$t=4$
A_1	askp ₅	askp ₄	askp ₃	tell (c_1)	
A_2	tell (c_1)				
A_3		tell (c_2)			

Fig. 12. A timeline for the execution of $A_1 \parallel A_2 \parallel A_3$.

9 Denotational Semantics for *tscpp-i*

In this section we define a denotational characterization of the operational semantics for *tscpp-i*. Differently from the denotational semantics for the maximal parallelism version presented in Section. 6.2, here for computational states we consider triples rather than pairs, as ω -actions have to be distinguished from τ -actions. This difference leads to a different technical development.

Our denotational model for *tscpp-i* associates with a process a set of timed reactive sequences of the form $\langle \sigma_1, \gamma_1, \xi_1 \rangle \cdots \langle \sigma_n, \gamma_n, \xi_n \rangle \langle \sigma, \sigma, \omega \rangle$. Any triple $\langle \sigma_i, \gamma_i, \xi_i \rangle$ represents a reaction (a computation step) of the given process at time i : intuitively, the process transforms the global store from σ_i to γ_i by performing a transition step labeled by ξ_i or, in other words, σ_i is the assumption on the external environment, ξ_i is the label of the performed step while γ_i is the contribution of the process itself (which entails always the assumption). The last pair denotes a “stuttering step”, in which the agent **Success** has been reached. In the following we will assume that each timed reactive sequence $\langle \sigma_1, \gamma_1, \xi_1 \rangle \cdots \langle \sigma_{n-1}, \gamma_{n-1}, \xi_{n-1} \rangle \langle \sigma_n, \sigma_n, \omega \rangle$ satisfies the following condition: $\gamma_i \vdash \sigma_i$ and $\sigma_j \vdash \gamma_{j-1}$, for any $i \in [1, n-1]$ and $j \in [2, n]$.

The basic idea underlying the denotational model then is that, differently from the operational semantics, inactive processes can always make a τ -step, where an inactive process is either suspended (due to the absence of the required constraint in the store) or it is a non-scheduled component of a parallel construct. These

additional τ -steps, which represent time-elapsing and are needed to obtain a compositional model in a simple way, are then added to denotations as triples of the form $\langle \sigma, \sigma, \tau \rangle$. For example, the denotation of the process $\mathbf{tell}(c) \rightarrow^a \mathbf{success}$ contains all the reactive sequences that have, as first element, a triple $\langle \sigma, \sigma \otimes c, \omega \rangle$ for any possible initial store σ with $(\sigma \otimes c) \Downarrow_{\emptyset} \not\prec a$, as these represent the action of adding the constraint c to the current store. However, such a denotation contains also sequences where the triple $\langle \sigma, \sigma \otimes c, \omega \rangle$ (still with $(\sigma \otimes c) \Downarrow_{\emptyset} \not\prec a$) is preceded by a finite sequence of triples of the form $\langle \sigma_1, \sigma_1, \tau \rangle \langle \sigma_2, \sigma_2, \tau \rangle \dots \langle \sigma_n, \sigma_n, \tau \rangle$. Such a sequence represents time-elapsing while the process is inactive because some other parallel process is scheduled.

The set of all reactive sequences for *tscpp-i* process is denoted by \mathcal{S}_i , its typical elements by $s, s_1 \dots$, while sets of reactive sequences are denoted by $S, S_1 \dots$ and ε indicates the empty reactive sequence. The operator \cdot denotes the operator that concatenates these sequences.

9.1 Compositionality of the Denotational Semantics for *tscpp-i* Processes

As in Section 6.2 for the *tscpp* version, we now introduce a denotational semantics $\mathcal{D}(F.A)(e)$ which is compositional by definition and where, for technical reasons, we represent explicitly the environment e which associates a denotation to each procedure identifier. More precisely, assuming that $Pvar$ denotes the set of procedure identifier, $Env_i = Pvar \rightarrow \mathcal{P}(\mathcal{S}_i)$, with typical element e , is the set of *environments*. Analogously to Section 6.2, given $e \in Env_i$, $p \in Pvar$ and $f \in \mathcal{P}(\mathcal{S}_i)$, we denote by $e' = e\{f/p\}$ the new environment such that $e'(p) = f$ and $e'(p') = e(p')$ for each procedure identifier $p' \neq p$.

Before defining formally the denotational semantics, we need to define the operators \bar{tell} , $\bar{\Sigma}$, \parallel , \bar{askp} and $\bar{\exists}x$, analogous to those given in Section 6.2 for the maximal parallelism language.

Definition 9 (Semantic operators for tscpp-i)

Let S, S_i be sets of reactive sequences, c, c_i be constraints. Moreover let \succ_i be either of the form \succ^{a_i} or \succ_{ϕ_i} , defined as in Definition 5. Then we define the operators \bar{tell} , $\bar{\Sigma}$, \parallel , \bar{askp} and $\bar{\exists}x$ as follows:

The (valued) tell operator $\bar{tell}^a : \mathcal{C} \times \wp(\mathcal{S}_i) \rightarrow \wp(\mathcal{S}_i)$ ($\bar{tell}_{\phi} : \mathcal{C} \times \wp(\mathcal{S}_i) \rightarrow \wp(\mathcal{S}_i)$) is the least function (w.r.t. the ordering induced by \subseteq) which satisfies the following equation

$$\begin{aligned} \bar{tell}^a(c, S) &= \{s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma \otimes c, \omega \rangle \cdot s', \sigma \otimes c \Downarrow_{\emptyset} \not\prec a \text{ and } s' \in S\} \cup \\ &\quad \{s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, \tau \rangle \cdot s' \text{ and } s' \in \bar{tell}^a(c, S)\}. \\ \bar{tell}_{\phi}(c, S) &= \{s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma \otimes c, \omega \rangle \cdot s', \sigma \otimes c \not\sqsubseteq \phi \text{ and } s' \in S\} \cup \\ &\quad \{s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, \tau \rangle \cdot s' \text{ and } s' \in \bar{tell}_{\phi}(c, S)\}. \end{aligned}$$

The guarded choice The semantic choice operator

$\bar{\sum}_{i=1}^n : (\mathcal{C} \times \wp(\mathcal{S}_i)) \times \cdots \times (\mathcal{C} \times \wp(\mathcal{S}_i)) \rightarrow \wp(\mathcal{S}_i)$ is the least function which satisfies the following equation:

$$\begin{aligned} \bar{\sum}_{i=1}^n c_i \succ_i S_i = & \{ s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, \omega \rangle \cdot s', \\ & \sigma \succ_h c_h \text{ and } s' \in S_h \text{ for an } h \in [1, n] \} \\ & \cup \\ & \{ s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, \tau \rangle \cdot s' \text{ and } s' \in \bar{\sum}_{i=1}^n c_i \succ_i S_i \}. \end{aligned}$$

Parallel Composition. Let $\bar{\parallel} \in \mathcal{S}_i \times \mathcal{S}_i \rightarrow \mathcal{S}_i$ be the (commutative and associative) partial operator defined by induction on the length of the sequences as follows:

$$\begin{aligned} \langle \sigma, \sigma, \omega \rangle \bar{\parallel} \langle \sigma, \sigma, \omega \rangle &= \langle \sigma, \sigma, \omega \rangle \\ \langle \sigma, \sigma', x \rangle \cdot s \bar{\parallel} \langle \sigma, \sigma, \tau \rangle \cdot s' &= \langle \sigma, \sigma, \tau \rangle \cdot s' \bar{\parallel} \langle \sigma, \sigma', x \rangle \cdot s = \langle \sigma, \sigma', x \rangle \cdot (s \bar{\parallel} s'), \end{aligned}$$

where $x \in \{\omega, \tau\}$.

We define the operator $S_1 \bar{\parallel} S_2$ on sets as the image of $\mathcal{S}_i \times \mathcal{S}_i$ under the above operator.

The (valued) askp operator $\text{askp}(t)^a : \mathcal{C} \times \wp(\mathcal{S}_i) \times \wp(\mathcal{S}_i) \rightarrow \wp(\mathcal{S}_i)$ ($\text{askp}(t)_\phi : \mathcal{C} \times \wp(\mathcal{S}_i) \times \wp(\mathcal{S}_i) \rightarrow \wp(\mathcal{S}_i)$), with $t > 0$, is defined as:

$$\begin{aligned} \text{askp}(t)^a(c, S_1, S_2) = & \{ s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, \omega \rangle \cdot s' \text{ and} \\ & \text{either } \sigma \succ^a c \text{ and } s \in S_1 \\ & \text{or } \sigma \Downarrow_\emptyset < a \text{ and } s \in S_2 \} \cup \\ & \{ s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, x \rangle \cdot s', s' \in \text{askp}(t-1)^a(c, S_1, S_2) \\ & \text{and either } x = \tau \\ & \text{or } x = \omega, \sigma \not\succ c \text{ and } \sigma \Downarrow_\emptyset \not\prec a \}. \\ \text{askp}(t)_\phi(c, S_1, S_2) = & \{ s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma', \omega \rangle \cdot s' \text{ and} \\ & \text{either } \sigma \succ_\phi c \text{ and } s \in S_1 \\ & \text{or } \sigma \sqsubset \phi \text{ and } s \in S_2 \} \cup \\ & \{ s \in \mathcal{S}_i \mid s = \langle \sigma, \sigma, x \rangle \cdot s', s' \in \text{askp}(t-1)_\phi(c, S_1, S_2) \\ & \text{and either } x = \tau \\ & \text{or } x = \omega, \sigma \not\succ c \text{ and } \sigma \not\sqsubset \phi \}. \end{aligned}$$

The (valued) askp operator $\text{askp}(0)^a : \mathcal{C} \times \wp(\mathcal{S}_i) \times \wp(\mathcal{S}_i) \rightarrow \wp(\mathcal{S}_i)$ ($\text{askp}(0)_\phi : \mathcal{C} \times \wp(\mathcal{S}_i) \times \wp(\mathcal{S}_i) \rightarrow \wp(\mathcal{S}_i)$) is the least function which satisfies the following equation

$$\begin{aligned} \text{askp}(0)^a(c, S_1, S_2) = & \{ s \in \mathcal{S}_i \mid \text{either } s = \langle \sigma, \sigma, \omega \rangle \cdot s' \text{ and } s' \in S_2 \\ & \text{or } s = \langle \sigma, \sigma, \tau \rangle \cdot s' \\ & \text{and } s' \in \text{askp}(0)^a(c, S_1, S_2) \}. \\ \text{askp}(0)_\phi(c, S_1, S_2) = & \{ s \in \mathcal{S}_i \mid \text{either } s = \langle \sigma, \sigma, \omega \rangle \cdot s' \text{ and } s' \in S_2 \\ & \text{or } s = \langle \sigma, \sigma, \tau \rangle \cdot s' \\ & \text{and } s' \in \text{askp}(0)_\phi(c, S_1, S_2) \}. \end{aligned}$$

The hiding operator The semantic hiding operator can be defined as follows:

$$\bar{\exists}xS = \{ s \in \mathcal{S}_i \mid \text{there exists } s' \in S \text{ such that } s = s'[x/y] \text{ with } y \text{ new} \}$$

where $s'[x/y]$ denotes the sequence obtained from s' by replacing the variable x for the variable y that we assume to be new.⁷

It is immediate to see that the previous semantic operators are well defined, that is, the least function which satisfies the equations actually exists and can be obtained by a standard fix-point construction. The $tell$, \sum , \parallel , $askp$ and $\exists x$ operators have the expected definition, including the mentioned addition of τ -steps.

In the semantic parallel operator (acting on sequences) we require that at each point of time at most one ω -action is present and the two arguments of the operator agree with respect to the contribution of the environment (the first component of the triple). We also require that the two arguments have the same length (in all other cases the parallel composition is assumed being undefined): this is necessary to reflect the passage of time since the i -th element of any sequence corresponds to the given processes action on the i -th time step. Even though we merge point-wise sequences of the same length, this models an interleaving approach for ω -actions, because of the previously mentioned addition of τ -steps to denotations. Concerning the semantic choice operator, a sequence in $\sum_{i=1}^n c_i \succ_i S_i$ consists of an initial period of waiting for a store which satisfies one of the guards. During this waiting period, only the environment is active by producing the constraint σ , while the process itself generates the stuttering steps $\langle \sigma, \sigma, \tau \rangle$. When the store is strong enough to satisfy a guard, that is to entail a c_h and to satisfy the condition on the cut level, then the resulting sequence is obtained by adding $s' \in S_h$ to the initial waiting period.

We can define the denotational semantics \mathcal{D} as follows. Here, $Process_i$ denotes the set of *tscpp*- i processes.

Definition 10 (Processes Semantics)

We define the semantics $\mathcal{D} \in Process_i \rightarrow \mathcal{P}(\mathcal{S}_i)$ is the least function with respect to the ordering induced by the set-inclusion, which satisfies the equations in Figure 13

Also \mathcal{D} is well defined and can be obtained by a fix-point construction. To see this, let us define an interpretation as a mapping $I : Process_i \rightarrow \wp(\mathcal{S}_i)$. Then let us denote by \mathcal{I} the cpo of all the interpretations (with the ordering induced by \subseteq). To the equations in Figure 13, we can then associate a monotonic (and continuous) mapping $\mathcal{F} : \mathcal{I} \rightarrow \mathcal{I}$ defined by the equations of Figure 13, provided that we replace the symbol \mathcal{D} for $\mathcal{F}(I)$, we delete the environment e and that we replace equation **F9** for the following one: $\mathcal{F}(I)(F.p(x)) = I(F.ask(\bar{1}) \rightarrow A)$.

Then, one can easily prove that a function satisfies the equations in Figure 13 iff it is a fix-point of the function \mathcal{F} . Because this function is continuous (on a cpo), well known results ensure us that its least fix-point exists and it equals \mathcal{F}^ω , where the powers are defined as follows: $\mathcal{F}^0 = I_0$ (this is the least interpretation which maps any process to the empty set); $\mathcal{F}^n = \mathcal{F}(\mathcal{F}^{n-1})$ and $\mathcal{F}^\omega = lub\{\mathcal{F}^n | n \geq 0\}$ (where lub is the least upper bound on the cpo \mathcal{I}).

⁷ As before, we assume that each time that we consider a new applications of the operator \exists we use a new, different y .

- F1** $\mathcal{D}(F.\text{success})(e) = \{\langle \sigma_1, \sigma_1, \tau \rangle \langle \sigma_2, \sigma_2, \tau \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle \in \mathcal{S}_i \mid n \geq 1\}$
- F2** $\mathcal{D}(F.\text{tell}(c) \rightarrow^a A)(e) = \bar{\text{tell}}^a(c, \mathcal{D}(F.A)(e))$
- F3** $\mathcal{D}(F.\text{tell}(c) \rightarrow_\phi A)(e) = \bar{\text{tell}}_\phi(c, \mathcal{D}(F.A)(e))$
- F4** $\mathcal{D}(F.\sum_{i=1}^n \text{ask}(c_i) \rightarrow_i A_i)(e) = \sum_{i=1}^n c_i \succ_i \mathcal{D}(F.A_i)(e)$
- F5** $\mathcal{D}(F.\text{askp}_t(c)?^a A:B)(e) = \bar{\text{askp}}(t)^a(c, \mathcal{D}(F.A)(e), \mathcal{D}(F.B)(e))$
- F6** $\mathcal{D}(F.\text{askp}_t(c)?_\phi A:B)(e) = \bar{\text{askp}}(t)_\phi(c, \mathcal{D}(F.A)(e), \mathcal{D}(F.B)(e))$
- F7** $\mathcal{D}(F.A \parallel B)(e) = \mathcal{D}(F.A)(e) \parallel \mathcal{D}(F.B)(e)$
- F8** $\mathcal{D}(F.\exists x A)(e) = \exists x \mathcal{D}(F.A)(e)$
- F9** $\mathcal{D}(F.p(x))(e) = \mu \Psi^i$ where $\Psi^i(f) = \mathcal{D}(F \setminus \{p\}.\text{ask}(\bar{1}) \rightarrow A)(e\{f/p\})$,
 $p(x) :: A \in F$

Fig. 13. The semantics $\mathcal{D}(F.A)(e)$ for *tscpp-i*.

9.2 Correctness of the Denotational Semantics for *tscpp-i* Processes

As for the correctness of the denotational semantics presented in Section 6.1, at each step, the assumption on the current store must be equal to the store produced by the previous step. In other words, for any two consecutive steps $\langle \sigma_i, \sigma'_i, x_i \rangle \langle \sigma_{i+1}, \sigma'_{i+1}, x_{i+1} \rangle$ we must have $\sigma'_i = \sigma_{i+1}$. Furthermore, triples containing τ -actions do not correspond to observable computational steps, as these involve ω -actions only.

Definition 11 (Connected Sequences in tscpp-i)

Let $s = \langle \sigma_1, \sigma'_1, x_1 \rangle \langle \sigma_2, \sigma'_2, x_2 \rangle \cdots \langle \sigma_n, \sigma'_n, \omega \rangle$ be a reactive sequence. We say that s is connected if $\sigma_1 = \bar{1}$, $\sigma_i = \sigma'_{i-1}$ and $x_j = \omega$ for each i, j , $2 \leq i \leq n$ and $1 \leq j \leq n-1$.

According to the previous definition, a sequence is connected if all the information assumed on the tuple space is produced by the process itself and only ω -actions are involved. To be defined as connected, a sequence must also have $\bar{1}$ as the initial constraint. A connected sequence represents a *tscpp-i* computation, as it will be proved in the remaining of this section.

In order to prove the correctness of the denotational semantics, we use a modified transition system T' , where inactive (either suspended or not scheduled) processes can perform τ -actions. When considering our notions of observables, we can prove that such a modified transition system is equivalent to the previous one and agrees with the denotational model.

The new transition system T' is obtained from the one in Figure 11 by deleting rule **Q6** and by adding the rules **Q0'**, **Q1'**, **Q2'**, **Q3'**, **Q4'**, **Q7'**, **Q8'**, **Q14'** and **Q19'**, contained in Figure 14. We denote by \Rightarrow the relation defined by T' .

The observables induced by the transition system T' are formally defined as follows.

Q0'	$\langle \text{success}, \sigma \rangle \xRightarrow{\tau} \langle \text{success}, \sigma \rangle$	success
Q1'	$\langle \text{tell}(c) \rightarrow^a A, \sigma \rangle \xRightarrow{\tau} \langle \text{tell}(c) \rightarrow^a A, \sigma \rangle$	V-TellI
Q2'	$\langle \text{tell}(c) \rightarrow_\phi A, \sigma \rangle \xRightarrow{\tau} \langle \text{tell}(c) \rightarrow_\phi A, \sigma \rangle$	Tell
Q3'	$\langle \text{ask}(c) \rightarrow^a A, \sigma \rangle \xRightarrow{\tau} \langle \text{ask}(c) \rightarrow^a A, \sigma \rangle$	V-ask
Q4'	$\langle \text{ask}(c) \rightarrow_\phi A, \sigma \rangle \xRightarrow{\tau} \langle \text{ask}(c) \rightarrow_\phi A, \sigma \rangle$	Ask
Q7'	$\langle \Sigma_{i=1}^n E_i, \sigma \rangle \xRightarrow{\tau} \langle \Sigma_{i=1}^n E_i, \sigma \rangle$	Nondet
Q8'	$\langle p(x), \sigma \rangle \xRightarrow{\tau} \langle A, \sigma \rangle \quad p(x) :: A \in F$	P-call
Q14'	$\langle \text{askp}_0(c)?^a A:B, \sigma \rangle \xRightarrow{\tau} \langle \text{askp}_0(c)?^a A:B, \sigma \rangle$	V-askp5
Q19'	$\langle \text{askp}_0(c)?_\phi A:B, \sigma \rangle \xRightarrow{\tau} \langle \text{askp}_0(c)?_\phi A:B, \sigma \rangle$	Askp5

Fig. 14. The τ -rules for *tscpp-i*.*Definition 12*

Let $P = F.A$ be a *tscpp-i* process. We define

$$\mathcal{O}_{io}^{i'}(P) = \{\gamma \Downarrow_{Fv(A)} \mid \langle A, \bar{\mathbf{I}} \rangle \xRightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle\},$$

where **Success** is any agent which contains only occurrences of the agent **success** and of the operator \parallel .

Lemma 3 shows that the modified transition system agrees with the original one when considering our notion of observables.

We first need some definitions and technical lemmata. In the following, given two agents A and B , we say that $A \simeq B$ if and only if B is obtained from A by replacing an agent of the form $\exists x A_1$ in A with $A_1[x/y]$, where y is new in A . \approx denotes the reflexive and transitive closure of \simeq . The following lemmata hold.

Lemma 1

Let $F.A$ and $F.B$ be *tscpp-i* processes such that $A \approx B$. Then for each store σ and for $x \in \{\omega, \tau\}$

$$\langle F.A, \sigma \rangle \xrightarrow{x} \langle F.C, \sigma' \rangle \text{ if and only if } \langle F.B, \sigma \rangle \xrightarrow{x} \langle F.C, \sigma' \rangle.$$

Proof

The proof is immediate, by using rule **Q9** and by a straightforward inductive argument. \square

From the above Lemma we derive the following corollary:

Corollary 1

Let $F.A$ and $F.B$ be *tscpp-i* processes such that $A \approx B$. Then for each store σ , $\langle F.A, \sigma \rangle \xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle$ if and only if $\langle F.B, \sigma \rangle \xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle$.

Lemma 2

Let $P = F.A$ be a *tscpp-i* process. Then for each store σ ,

1. $\langle F.A, \sigma \rangle \xRightarrow{\tau} \langle F.B, \sigma' \rangle$ if and only if $\sigma = \sigma'$ and
 - either $\langle F.A, \sigma \rangle \xrightarrow{\tau} \langle F.C, \sigma \rangle$ and $C \approx B$
 - or $\langle F.A, \sigma \rangle \not\xrightarrow{\tau}$ and $B \approx A$.
2. $\langle F.A, \sigma \rangle \xRightarrow{\omega} \langle F.B, \sigma' \rangle$ if and only if $\langle F.A, \sigma \rangle \xrightarrow{\omega} \langle F.C, \sigma' \rangle$ and $C \approx B$.

Proof

1. The proof is by induction on the complexity of the agent A .
 - A is of the form **success**, **tell**(c) $\rightarrow^a A$, **tell**(c) $\rightarrow_\phi A$, **ask**(c) $\rightarrow^a A$, **ask**(c) $\rightarrow_\phi A$, $\Sigma_{i=1}^n E_i$, $p(x)$, **askp**₀(c) $?^a A:B$ and **askp**₀(c) $?_\phi A:B$.
 The proof is immediate by observing that by the rules in Figure 11, $\langle A, \sigma \rangle \not\xrightarrow{\tau}$ and by the rules in Figure 14, $\langle A, \sigma \rangle \xRightarrow{\tau} \langle B, \sigma' \rangle$ if and only if $\langle A, \sigma \rangle = \langle B, \sigma' \rangle$.
 - A is of the form **askp** _{t} (c) $?^a A_1:A_2$ (**askp** _{t} (c) $?_\phi A_1:A_2$), with $t > 0$.
 The proof is immediate since both the transition systems use the rule **Q13** (**Q18**) of Figure 11.
 - If A is of the form $A_1 \parallel A_2$.
 In this case, by definition of the transition system T' and by using rule **Q5** of Figure 11, for each store σ ,

$$\langle A_1 \parallel A_2, \sigma \rangle \xRightarrow{\tau} \langle B_1 \parallel B_2, \sigma' \rangle \text{ if and only if } \langle A_1, \sigma \rangle \xRightarrow{\tau} \langle B_1, \sigma' \rangle \text{ and } \langle A_2, \sigma \rangle \xRightarrow{\tau} \langle B_2, \sigma' \rangle$$
 (the symmetric case is analogous and hence it is omitted).
 By inductive hypothesis this holds if and only if $\sigma' = \sigma$ and for $i = 1, 2$
 - either $\langle A_i, \sigma \rangle \xrightarrow{\tau} \langle C_i, \sigma \rangle$ and $C_i \approx B_i$
 - or $\langle A_i, \sigma \rangle \not\xrightarrow{\tau}$ and $B_i \approx A_i$.
 If there exists $i \in [1, 2]$ such that $\langle A_i, \sigma \rangle \xrightarrow{\tau} \langle C_i, \sigma \rangle$ then the thesis follows by using either rule **Q5** or rule **Q6**.
 Otherwise $\langle A_1 \parallel A_2, \sigma \rangle \not\xrightarrow{\tau}$. Then the thesis follows since by the previous results $B_1 \parallel B_2 \approx A_1 \parallel A_2$.
 - A is of the form $\exists x A_1$. By rule **Q9** of Figure 11 for each store σ ,

$$\langle \exists x A_1, \sigma \rangle \xRightarrow{\tau} \langle B, \sigma' \rangle \text{ if and only if } \langle A_1[x/y], \sigma \rangle \xRightarrow{\tau} \langle B, \sigma' \rangle$$
 By inductive hypothesis this holds if and only if $\sigma' = \sigma$
 - either $\langle A_1[x/y], \sigma \rangle \xrightarrow{\tau} \langle C, \sigma \rangle$ and $C \approx B$
 - or $\langle A_1[x/y], \sigma \rangle \not\xrightarrow{\tau}$ and $B \approx A_1[x/y]$.
 Therefore, by using rule **Q9** of Figure 11 and since $\exists x A_1 \approx A_1[x/y]$, we have that
 - either $\langle \exists x A_1, \sigma \rangle \xrightarrow{\tau} \langle C, \sigma \rangle$ and $C \approx B$
 - or $\langle \exists x A_1, \sigma \rangle \not\xrightarrow{\tau}$ and $B \approx \exists x A_1$ and then the thesis.
2. The proof is analogous to the previous one and hence it is omitted.

□

Lemma 3

Let $P = F.A$ be a *tscpp-i* process. Then $\mathcal{O}_{io}^{i'}(P) = \mathcal{O}_{io}^i(P)$.

Proof

We prove that there exists a computation $\langle A, \sigma \rangle \xRightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle$ if and only if there exists a computation $\langle A, \sigma \rangle \xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle$. Then the thesis follows by definition of $\mathcal{O}_{io}^i(P)$ and $\mathcal{O}_{io}^{i'}(P)$. The proof is by induction on the length of the computation $\langle A, \sigma \rangle \xRightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle$.

$n = 1$) In this case $A = \mathbf{Success}$ and then the thesis.

$n > 1$) In this case

$$\begin{aligned}
 \langle A, \sigma \rangle &\xRightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle && \text{iff} \\
 &(\text{by definition}) \\
 \langle A, \sigma \rangle &\xRightarrow{\omega} \langle A_1, \sigma_1 \rangle \text{ and } \langle A_1, \sigma_1 \rangle \xRightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle && \text{iff} \\
 &(\text{by inductive hypothesis}) \\
 \langle A, \sigma \rangle &\xRightarrow{\omega} \langle A_1, \sigma_1 \rangle \text{ and } \langle A_1, \sigma_1 \rangle \xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle && \text{iff} \\
 &(\text{by Point 2 of Lemma 2}) \\
 \langle A, \sigma \rangle &\xrightarrow{\omega} \langle A_2, \sigma_1 \rangle, A_2 \approx A_1 \text{ and } \langle A_1, \sigma_1 \rangle \xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle && \text{iff} \\
 &(\text{by Corollary 1}) \\
 \langle A, \sigma \rangle &\xrightarrow{\omega} \langle A_2, \sigma_1 \rangle \text{ and } \langle A_2, \sigma_1 \rangle \xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle && \text{iff} \\
 &(\text{by definition}) \\
 \langle A, \sigma \rangle &\xrightarrow{\omega}^* \langle \mathbf{Success}, \gamma \rangle.
 \end{aligned}$$

□

We can now easily prove that, given our definition of \mathcal{D} , the modified transition system T' agrees with the denotational model.

Theorem 2

For any *tscpp-i* process $P = F.A$ we have

$$\begin{aligned}
 \mathcal{O}_{io}^{i'}(P) = \{ \sigma_n \Downarrow_{Fv(A)} \mid & \text{there exists a connected sequence } s \in \mathcal{D}(P) \text{ such that} \\
 s = \langle \sigma_1, \sigma_2, \omega \rangle \langle \sigma_2, \sigma_3, \omega \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle \}. &
 \end{aligned}$$

Proof

We prove by induction on the complexity of the agent A that

$$\begin{aligned}
 \mathcal{D}(P) = \{ s \mid & s = \langle \sigma_1, \sigma'_1, x_1 \rangle \langle \sigma_2, \sigma'_2, x_2 \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle, A_1 = A, \\
 & \text{for } i \in [1, n-1], \langle A_i, \sigma_i \rangle \xrightarrow{x_i} \langle A_{i+1}, \sigma'_i \rangle \text{ and } A_n = \mathbf{Success} \}.
 \end{aligned}$$

Then the proof follows by definition of $\mathcal{O}_{io}^{i'}(P)$.

When the *tscpp-i* P is not of the form $F.B \parallel C$ the thesis follows immediately from the close correspondence between the rules of the transition system and the definition of the denotational semantics.

Assume now that P is of the form $F.B \parallel C$. By definition of the denotational

semantics, $s \in \mathcal{D}(P)$ if and only if $s = \langle \sigma_1, \sigma'_1, x_1 \rangle \langle \sigma_2, \sigma'_2, x_2 \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle$ and there exist $s' \in \mathcal{D}(F.B)$ and $s'' \in \mathcal{D}(F.C)$,

$$\begin{aligned} s' &= \langle \sigma_1, \kappa'_1, x'_1 \rangle \langle \sigma_2, \kappa'_2, x'_2 \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle \quad \text{and} \\ s'' &= \langle \sigma_1, \kappa''_1, x''_1 \rangle \langle \sigma_2, \kappa''_2, x''_2 \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle, \end{aligned}$$

such that for each $i \in [1, n-1]$,

$$\begin{aligned} x_i = \tau \text{ if and only if } & x'_i = x''_i = \tau \text{ and in this case } \sigma'_i = \kappa'_i = \kappa''_i = \sigma_i, \\ x_i = \omega \text{ if and only if } & \text{either } x'_i = \omega, x''_i = \tau, \kappa'_i = \sigma'_i \text{ and } \kappa''_i = \sigma_i \\ & \text{or } x'_i = \tau, x''_i = \omega, \kappa'_i = \sigma'_i \text{ and } \kappa''_i = \sigma_i. \end{aligned} \quad (2)$$

By inductive hypothesis $s' \in \mathcal{D}(F.B)$ and $s'' \in \mathcal{D}(F.C)$ if and only if

$$\begin{aligned} \langle B_i, \sigma_i \rangle &\xrightarrow{x'_i} \langle B_{i+1}, \kappa'_i \rangle \text{ for } i \in [1, n-1], \quad B_1 = B \text{ and } B_n = \mathbf{Success}, \\ \langle C_i, \sigma_i \rangle &\xrightarrow{x''_i} \langle C_{i+1}, \kappa''_i \rangle \text{ for } i \in [1, n-1], \quad C_1 = C \text{ and } C_n = \mathbf{Success}. \end{aligned} \quad (3)$$

Therefore, by Rule **R8** and by (2), we have that (3) holds if and only if

$$\begin{aligned} \langle B_i \parallel C_i, \sigma_i \rangle &\xrightarrow{x_i} \langle B_{i+1} \parallel C_{i+1}, \sigma_{i+1} \rangle \text{ for } i \in [1, n-1], \\ B_1 \parallel C_1 &= B \parallel C \text{ and } B_n \parallel C_n = \mathbf{Success} \end{aligned}$$

and then the thesis. \square

Thus we obtain the following correctness result whose proof is immediate from the previous theorems.

Corollary 2 (Correctness of tscpp-i)

For any *tscpp-i* process $P = F.A$ we have

$$\begin{aligned} \mathcal{O}_{io}^i(P) = \{ \sigma_n \Downarrow_{Fv(A)} \mid & \text{there exists a connected sequence } s \in \mathcal{D}(P) \text{ such that} \\ & s = \langle \sigma_1, \sigma_2, \omega \rangle \langle \sigma_2, \sigma_3, \omega \rangle \cdots \langle \sigma_n, \sigma_n, \omega \rangle \}. \end{aligned}$$

10 Related Work

By comparing this work with other timed languages using crisp constraints (instead of soft ones as in this paper) as (Saraswat et al. 1996; Saraswat et al. 1994), there are three main differences we can find out.

First, the computational model of both the languages *tcc* (Saraswat et al. 1994) and *default tcc* (Saraswat et al. 1996) is inspired by that one of synchronous languages: each time interval is identified with the time needed for a *ccp* process to terminate a computation. Clearly, in order to ensure that the next time instant is reached, the (default) *ccp* program has to be always terminating; thus, it is assumed that it does not contain recursion. On the other hand, we directly introduce a timed interpretation of the usual programming constructs of *ccp* by considering the primitive *ccp* constructs **ask** and **tell** as the elementary actions whose evaluation takes one time-unit. Therefore, in our model, each time interval is identified with the time needed for the underlying constraint system to accumulate the tells and to answer the queries (asks) issued at each computation step by the processes of the system. For the definition of our *tscpp* agents we do not need any restriction on recursion to

ensure that the next time instant is reached, since at each moment there are only a finite number of parallel agents, and the next moment in time occurs as soon as the underlying constraint system has responded to the initial actions of all the current agents of the system.

A second difference relies in the transfer of information across time boundaries. In (Saraswat et al. 1994) and (Saraswat et al. 1996), the programmer has to explicitly transfer the (positive) information from a time instant to the next one, by using special primitives that allow one to control the temporal evolution of the system. In fact, at the end of a time interval all the constraints accumulated and all the processes suspended are discarded, unless they are arguments to a specific primitive. On the contrary, no explicit transfer is needed in *tsccp*, since the computational model is based on the monotonic evolution of the store which is usual in *ccp*.

A third relevant difference is in (Saraswat et al. 1994) and (Saraswat et al. 1996) the authors present deterministic languages while our language allows for non-determinism. These three differences also hold between (Saraswat et al. 1994) or (Saraswat et al. 1996), and the original crisp version of the language, i.e., *tccp* (de Boer et al. 2000).

In (Olate et al. 2007), the authors generalize the model in (Saraswat et al. 1994) in order to extend it with temporary parametric **ask** operations. Intuitively, these operations behave as persistent parametric asks during a time-interval, but may disappear afterwards. The presented extension goes in the direction of better modeling *mobile systems* with the use of private channels between the agents. However, also the agents in (Olate et al. 2007) show a deterministic behavior, instead of our not-deterministic choice.

Other timed extension of concurrent constraint programming have been proposed in (Nielsen and Valencia 2002; Palamidessi and Valencia 2001), however these languages, differently from *tsccp*, do not take into account quantitative aspects; therefore, this achievement represents a very important expressivity improvement with respect to related works. These have been considered by Di Pierro and Wiklicky, who have extensively studied probabilistic *ccp* (see for example (Di Pierro and Wiklicky 1998)). This language provides a construct for probabilistic choice which allows one to express randomness in a program, without assuming any additional structure on the underlying constraint system. This approach is therefore deeply different from ours. More recently, stochastic *ccp* has been introduced in (Bortolussi 2006) to model biological systems. This language is obtained by adding a stochastic duration to the ask and tell primitives, thus it differs from our solutions.

In literature we can find other proposals that are related to tuple-based kernel-languages instead of a constraint store, as *KLAIM* (de Nicola et al. 1998) (*A Kernel Language for Agents Interaction and Mobility*) or *SCEL* (De Nicola et al. 2011) (*Software Component Ensemble Language*) for instance. These languages are designed to study different properties of systems, as mobility and autonomicity of modeled agents. Their basic specification do not encompass time-based primitives, while mobility features are not present in any of the constraint-based languages reported in this section. The purpose of our language is to model systems where a level of preference and time-sensitive primitives (as a timeout) is required: a good

example is represented by agents participating to an auction, as the example given in Section 5.1.

In general, since semiring-based soft constraints allow one to express several quantitative features, our proposal provides a framework which can be instantiated to obtain a variety of specific extensions of *ccp*.

11 Conclusion and Future Work

We have presented the *tscpp* and *tscpp-i* in order to join together the expressive capabilities of soft constraints and timing mechanisms in a new programming framework. The agents modeled with these languages are able to deal with time and preference-dependent decisions that are often found during complex interactions. An application scenario can be represented by different entities that need to negotiate generic resources or services, as, for instance, during an auction process. Mechanisms as timeout and interrupt may model the wait for pending conditions or the triggering of some new events. All the *tscpp* and *tscpp-i* rules have been formally described by a transition system and, then, also with a denotational characterization of the operational semantics obtained with the use of *timed reactive sequences*. The resulting semantics has been proved to be compositional and correct.

About future work, a first improvement of the presented languages can be the inclusion of a *fail* agent in the syntax given in Definition 1 and Definition 7, and a semantics for the transition rules that lead to a failed computation, in case the guard on the transition rule cannot be enforced due to the preference of the store. In fact, the transition systems we have defined consider only successful computations. If this could be a reasonable choice in a don't know interpretation of the language it will lead to an insufficient analysis of the behavior in a pessimistic interpretation of the indeterminism.

At last, we would like to consider other time management strategies (as the one proposed in (Valencia 2003)), and to study how timing and non-monotonic constructs (Bistarelli and Santini 2011) can be integrated together.

References

- BERRY, G. AND GONTHIER, G. 1992. The ESTEREL synchronous programming language: design, semantics, implementation. *Sci. Comput. Program.* 19, 2, 87–152.
- BISTARELLI, S. 2004. *Semirings for Soft Constraint Solving and Programming (LNCS)*. Springer Verlag, London, UK.
- BISTARELLI, S., GABBRIELLI, M., MEO, M. C., AND SANTINI, F. 2008. Timed soft concurrent constraint programs. In *Coordination Models and Languages, 10th International Conference, COORDINATION*. Lecture Notes in Computer Science, vol. 5052. Springer, London, UK, 50–66.
- BISTARELLI, S., MONTANARI, U., AND ROSSI, F. 1997. Semiring-based constraint satisfaction and optimization. *J. ACM* 44, 2, 201–236.
- BISTARELLI, S., MONTANARI, U., AND ROSSI, F. 2006. Soft concurrent constraint programming. *ACM Trans. Comput. Logic* 7, 3, 563–589.
- BISTARELLI, S. AND SANTINI, F. 2011. A nonmonotonic soft concurrent constraint language to model the negotiation process. *Fundam. Inform.* 111, 3, 257–279.

- BORTOLUSSI, L. 2006. Stochastic concurrent constraint programming. *Electr. Notes Theor. Comput. Sci.* 164, 3, 65–80.
- BROOKES, S. D. 1993. Full abstraction for a shared variable parallel language. In *LICS*. IEEE Computer Society, Los Alamitos, CA, USA, 98–109.
- BUSI, N., GORRIERI, R., AND ZAVATTARO, G. 2000. Process calculi for coordination: From Linda to JavaSpaces. In *AMAST*. Springer-Verlag, London, UK, 198–212.
- DE BOER, F. S., GABBRIELLI, M., AND MEO, M. C. 2000. A timed concurrent constraint language. *Inf. Comput.* 161, 1, 45–83.
- DE BOER, F. S., GABBRIELLI, M., AND MEO, M. C. 2004. A timed Linda language and its denotational semantics. *Fundam. Inf.* 63, 4, 309–330.
- DE BOER, F. S. AND PALAMIDESSI, C. 1991. A fully abstract model for concurrent constraint programming. In *TAPSOFT '91: Proceedings of CAAP '91*. Vol. 1. Springer-Verlag New York, Inc., New York, USA, 296–319.
- DE NICOLA, R., FERRARI, G. L., LORETI, M., AND PUGLIESE, R. 2011. A language-based approach to autonomic computing. In *FMCO*, B. Beckert, F. Damiani, F. S. de Boer, and M. M. Bonsangue, Eds. Lecture Notes in Computer Science, vol. 7542. Springer, 25–48.
- DE NICOLA, R., FERRARI, G. L., AND PUGLIESE, R. 1998. KLAIM: A Kernel Language for Agents Interaction and Mobility. *IEEE Transactions on Software Engineering* 24, 5, 315–330.
- DI PIERRO, A. AND WIKLICKY, H. 1998. Probabilistic concurrent constraint programming: Towards a fully abstract model. In *MFCS '98: Proceedings of Mathematical Foundations of Computer Science*. Springer-Verlag, London, UK, 446–455.
- GALLIEN, J. AND GUPTA, S. 2007. Temporary and permanent buyout prices in online auctions. *Management Science* 53, 5, 814–833.
- HALBWACHS, N., CASPI, P., RAYMOND, P., AND PILAUD, D. 1991. The synchronous data-flow programming language LUSTRE. *Proceedings of the IEEE* 79, 9, 1305–1320.
- HAREL, D. 1987. Statecharts: A visual formalism for complex systems. *Sci. Comput. Program.* 8, 3, 231–274.
- JONSSON, B. 1985. A model and proof system for asynchronous networks. In *PODC '85: Proceedings ACM symposium on Principles of distributed computing*. ACM Press, New York, USA, 49–58.
- LE GUERNIC, P., LE BORGNE, M., GAUTIER, T., AND LE MAIRE, C. 1991. Programming real-time applications with signal. *Proceedings of the IEEE* 79, 9, 1321–1336.
- NIELSEN, M. AND VALENCIA, F. D. 2002. Temporal concurrent constraint programming: Applications and behavior. In *Formal and Natural Computing - Essays Dedicated to Grzegorz Rozenberg*. Springer-Verlag, London, UK, 298–324.
- OLARTE, C., PALAMIDESSI, C., AND VALENCIA, F. 2007. Universal timed concurrent constraint programming. In *Logic Programming, 23rd International Conference, ICLP*. LNCS, vol. 4670. Springer, London, UK, 464–465.
- PALAMIDESSI, C. AND VALENCIA, F. D. 2001. A temporal concurrent constraint programming calculus. In *CP '01: Proceedings of Principles and Practice of Constraint Programming*. Springer-Verlag, London, UK, 302–316.
- SARASWAT, V. 1989. Concurrent constraint programming languages. PhD thesis.
- SARASWAT, V., JAGADEESAN, R., AND GUPTA, V. 1996. Timed default concurrent constraint programming. *J. Symb. Comput.* 22, 5-6, 475–520.
- SARASWAT, V. AND RINARD, M. 1990. Concurrent constraint programming. In *POPL '90: Proceedings of Principles of programming languages*. ACM Press, New York, USA, 232–245.

- SARASWAT, V. A., JAGADEESAN, R., AND GUPTA, V. 1994. Foundations of timed concurrent constraint programming. In *Proceedings, Ninth Annual IEEE Symposium on Logic in Computer Science, LICS*. IEEE Computer Society, Los Alamitos, CA, USA, 71–80.
- VALENCIA, F. D. 2003. Timed concurrent constraint programming: Decidability results and their application to LTL. In *ICLP*. LNCS, vol. 2916. Springer, London, UK, 422–437.